

How to compare different loss functions and their risks

Ingo Steinwart
Modeling, Algorithms and Informatics Group, CCS-3
Los Alamos National Laboratory
`{ingo}@lanl.gov`

January 20, 2006

Abstract

Many learning problems are described by a risk functional which in turn is defined by a loss function, and a straightforward and widely-known approach to learn such problems is to minimize a (modified) empirical version of this risk functional. However, in many cases this approach suffers from substantial problems such as computational requirements in classification or robustness concerns in regression. In order to resolve these issues many successful learning algorithms try to minimize a (modified) empirical risk of a surrogate loss function, instead. Of course, such a surrogate loss must be “reasonably related” to the original loss function since otherwise this approach cannot work well. For classification good surrogate loss functions have been recently identified, and the relationship between the excess classification risk and the excess risk of these surrogate loss functions has been exactly described. However, beyond the classification problem little is known on good surrogate loss functions up to now. In this work we establish a general theory that provides powerful tools for comparing excess risks of different loss functions. We then apply this theory to several learning problems including (weighted) classification, regression, density estimation, and density level detection.

1 Introduction

In many machine learning problems the learning goal is described by a loss function and its associated risk. A typical example of such a learning problem is *binary classification*, where a training set $T := ((x_1, y_1), \dots, (x_n, y_n))$ is given and the goal is to predict the label $y \in Y := \{-1, 1\}$ for a new, unseen input sample $x \in X$. Commonly, it is assumed that the samples are drawn in an i.i.d. fashion from an unknown probability measure P on $X \times Y$. The simplest learning goal is then to have an almost minimal prediction error on average future data, i.e. to find a function $f : X \rightarrow \mathbb{R}$ such that

$$P\left(\{(x, y) \in X \times Y : \text{sign } f(x) \neq y\}\right)$$

is as small as possible. It is well-known and obvious that writing $L(y, t) := 1$ if $y \text{sign } t < 0$ and $L(y, t) := 0$ otherwise, the above probability equals the so-called classification risk

$$\mathcal{R}_{L,P}(f) := \int_{X \times Y} L(y, f(x)) dP(x, y)$$

and consequently, the learning goal is then to find a function f that (approximately) minimizes this risk. This form of a learning goal appears in many other learning problems, too. Probably the best known example of such a learning problem is real-valued *regression*, however there are many other

important learning problems like *weighted binary classification*, *multi-class classification*, *density estimation*, or *density level detection* which can also be described by a risk functional.

Having a learning problem of the above form a simple and well-known learning approach is the empirical risk minimization (ERM) method. Basically, the idea of ERM is to replace the unknown true risk $\mathcal{R}_{L,P}(\cdot)$ by its empirical counterpart based on the training set T and minimize this empirical L -risk over a suitable function class. Though this approach has a lot of theoretical merits, for classification it typically leads to NP-hard optimization problems (see e.g. [13]), and thus it is not computationally realizable. One way to resolve this issue is to use a *surrogate loss function* L_2 , e.g. the hinge loss as in support vector machines, and then to minimize the (perhaps even further modified) empirical L_2 -risk. Let us assume for a moment that such a learning method \mathcal{L} learns the surrogate learning problem defined by the loss L_2 , i.e. we know $\mathcal{R}_{L_2,P}(f_T) \rightarrow \mathcal{R}_{L_2,P}^*$, where f_T is the function produced by the learning method \mathcal{L} and $\mathcal{R}_{L_2,P}^* := \inf\{f : X \rightarrow \mathbb{R} \mid f \text{ measurable}\}$ is the smallest possible L_2 -risk. Then first question which then naturally arise is:

Question 1. *Does the convergence $\mathcal{R}_{L_2,P}(f_T) \rightarrow \mathcal{R}_{L_2,P}^*$ imply the convergence $\mathcal{R}_{L,P}(f_T) \rightarrow \mathcal{R}_{L,P}^*$?*

If we can find a positive answer to this question we know at least asymptotically that by solving the surrogate learning problem we actually also solve the learning problem we are interested in, and in this sense \mathcal{L} is a suitable learning approach. However, in many situations we are not only interested in such an asymptotic relation, but also in a more quantitative statement. This leads to the second question:

Question 2. *Does there exists an increasing function $\Upsilon : [0, \infty) \rightarrow [0, \infty)$ with $\Upsilon(0) = 0$, which is continuous in 0 and satisfies*

$$\mathcal{R}_{L,P}(f) - \mathcal{R}_{L,P}^* \leq \Upsilon(\mathcal{R}_{L_2,P}(f) - \mathcal{R}_{L_2,P}^*) ?$$

For the binary classification problem both questions have been intensively investigated in recent years (see e.g. [17, 18, 36, 26, 1]), and for margin-based loss functions L_2 , i.e. for loss functions of the form $L_2(y, t) = \varphi(yt)$, complete answers were established in [1]. However, using a surrogate loss function is a strategy that is not only interesting for binary classification. Indeed, this approach also has its benefits in the following learning problems (see Section 3 for a rigorous introduction):

Weighted binary classification. In weighted binary classification the two different types of errors are penalized differently. Consequently, the optimization problem of a simple ERM approach remains NP-hard to solve, and thus suitable surrogate loss functions are desired for a computational treatment. Despite the practical importance of this learning problem only very little is known in view of the above questions (see [19] for a conjecture related to the above questions).

Regression. The classical loss function for regression problems is the least squares loss. Though this loss is mathematically rather easy to handle and the corresponding learning algorithms are often computational feasible, it is well-known that minimizing an empirical risk based on the least squares loss is a method which is quite sensitive to outliers. Therefore, a number of more robust surrogate loss functions have been proposed in the last decades (see e.g. [14, 31, 24, 9]). Up to now, the analysis of such regression loss functions is typically conducted in a maximum-likelihood fashion which relates the loss functions to density models of the marginal distribution $P(Y|x)$. However, the maximum-likelihood approach requires knowledge about the marginal distribution which is usually not available and in addition, it does not give information regarding our two basic questions.

Density level detection. The common performance measure of the density level detection problem is a risk functional whose loss function is defined in terms of the *unknown* density (see e.g. [11, 20, 21, 30]). Consequently, the loss function is unknown and cannot be used to *a)* compute a test error, *b)* build a (modified) ERM approach, and *c)* compare different solutions to the density level detection problem. Therefore, having a reasonable surrogate risk is essential for dealing with the density level detection problem.

Density estimation. The common performance measures for the density estimation problem are p -norms of the difference between the unknown density and its estimate. Again, these performance measures are risks with respect to a loss function defined by the unknown density, and consequently, the density estimation problem suffers from the same problems as the density level detection problem does.

Obviously, the above learning problems are rather different regarding their defining risk functionals, and in addition, the above list is by now means complete (see e.g. the recent work on multi-class classification in [35, 29]). In order to systematically investigate our two questions it is therefore reasonable to establish a general theory first, and then apply this theory to the learning problem at hand. This is exactly the path we follow in this work. More precisely, we first establish a very general theory that describes how to relate the excess risks of different loss functions. In order to demonstrate this power of this theory, we then apply it to the learning problems described above. Let us briefly summarize our findings:

- **Weighted classification.** We show that a natural weighting method for margin-based loss functions allows us to translate the results of [1] to analogous results for weighted classification. In particular, we give positive answers to Question 1 and 2. Moreover, we show that the natural weighting method is the only one that allows such positive answers.
- **Regression.** We first show that the least squares loss is essentially the only loss function which can be used to find the regression function $\mathbb{E}(Y|x)$. For some large classes of *symmetric* noise distributions $P(Y|x)$ we then characterize the loss functions of the form $L_2(y, t) = \psi(y - t)$ which allow positive answers to the above questions if the target risk is the 1-norm distance between the regression function $\mathbb{E}(Y|.)$ and its prediction $f(.)$. Here it will turn out that the convexity and related, stronger notions such as strict convexity and uniform convexity play a crucial role. Moreover, we show that for unbounded, symmetric, but otherwise unspecified noise every loss function admitting a positive answer to Question 2 must grow at least as fast as the squared loss. Consequently, every ERM approach to find the regression function $\mathbb{E}(Y|x)$ either learns the problem only in a weak sense or is sensitive to outliers. Finally, we discuss in which cases and in which sense approximate risk minimizers f_T approximate the exact risk minimizer.
- **Density level detection.** It was recently shown in [27] that the density level detection problem can be solved by learning a binary classification problem, and in fact, the widely known excess mass approach (see e.g. [11, 20, 21, 30]) implicitly implements ERM of this classification problem. We extend the considerations of [27] and show that there exists no surrogate loss function which allows a positive answer to the second question without making assumptions on the density.
- **Density estimation.** A well-known heuristic (see e.g. [12, Chap. 14.2.4]) casts the density estimation problem into a supervised learning problem which is related to the supervised surrogate for the density level detection problem. However, to our best knowledge, there is

nothing known about this heuristic in view of the above questions. For convex loss functions we are able to give a weak positive answer to the first question. Furthermore, we present a result suggesting that the second answer cannot be positively answered.

The rest of this work is organized as follows: In Section 2 we develop the general theory on excess risks. Section 3 then contains the results for the above learning problems. The proofs of the results of Section 2 can be found in Section 4. In addition, this section contains a powerful tool ensuring the existence of measurable selections which is used to deal with the notoriously unpleasant measurability questions related to minimal risks. Finally, the appendix contains some additional information on stronger notions of convexity.

2 Definitions and Main Results

In this section we introduce the central definitions and present the main results of our general theory. In the first subsection general cost functions of a form similar to those considered in [32] are studied. In particular, the fundamental results that compare excess risks of different cost functions are established. In the second subsection these results are applied to loss functions of the form used in the introduction. Finally, in the third subsection we generalize the notion of loss functions to establish powerful tools for treating more advanced questions on loss functions and their risks.

2.1 The General Theory

Let us begin with some notations which will be used throughout this work. To this end let X be a nonempty set equipped with a σ -algebra \mathcal{X} . Recall that the completion $\hat{\mathcal{X}}$ of \mathcal{X} is defined by

$$\hat{\mathcal{X}} := \left\{ A \subset X : A \in X_\mu \text{ for all finite measures } \mu \text{ on } \mathcal{X} \right\},$$

where X_μ denotes the completion with respect to the measure μ . Furthermore, the measurable space (X, \mathcal{X}) is called *complete* if $\mathcal{X} = \hat{\mathcal{X}}$. Throughout this work we assume that (X, \mathcal{X}) is complete. Moreover, Y always denotes Polish space, i.e. a topological space whose topology can be described by a complete and separable metric, and we always equip Y with its Borel σ -algebra. In particular recall that all open or closed subsets of \mathbb{R} are Polish spaces.

For a probability measure P on $X \times Y$ we denote the marginal distribution on X by P_X . Furthermore, $P(\cdot | \cdot) : \mathcal{Y} \times X \rightarrow [0, 1]$ stands for a *fixed* regular conditional probability, so that we have

$$\int_{X \times Y} f dP = \int_X \int_Y f(x, y) dP(y|x) dP_X(x)$$

for all measurable functions $f : X \times Y \rightarrow [0, \infty]$. Finally, throughout this subsection \mathcal{A} denotes a non-empty but otherwise arbitrary set. We begin with some fundamental definitions.

Definition 2.1 (Cost function) *A function $L : X \times Y \times \mathcal{A} \rightarrow [0, \infty]$ is called a cost function if $L(\cdot, \cdot, \alpha) : X \times Y \rightarrow [0, \infty]$ is measurable for all $\alpha \in \mathcal{A}$*

Our main research objects will be the risks that are associated to cost functions via a integration procedure. The following definition introduces such risks:

Definition 2.2 (Risk) *Let $L : X \times Y \times \mathcal{A} \rightarrow [0, \infty]$ be a cost function and P be a distribution on $X \times Y$. We define the L -risk of $\alpha \in \mathcal{A}$ by*

$$\mathcal{R}_{L,P}(\alpha) := \int_{X \times Y} L(x, y, \alpha) dP(x, y) = \int_X \int_Y L(x, y, \alpha) dP(y|x) dP_X(x). \quad (1)$$

Furthermore, the minimal L -risk is denoted by $\mathcal{R}_{L,P}^* := \inf_{\alpha \in \mathcal{A}} \mathcal{R}_{L,P}(\alpha)$.

Using the regular conditional probability $P(y|x)$, the risk $\mathcal{R}_{L,P}(\alpha)$ can be computed by an iterated integral as we have seen in (1). Since the inner integral in (1) will be of fundamental importance in our analysis we introduce:

Definition 2.3 (Inner risk) Let $L : X \times Y \times \mathcal{A} \rightarrow [0, \infty]$ be a cost function and Q be a distribution on Y . We define the inner L -risk of an element $\alpha \in \mathcal{A}$ by

$$\mathcal{C}_{L,Q,x}(\alpha) := \int_Y L(x, y, \alpha) dQ(y) \quad x \in X.$$

Furthermore, the minimal inner L -risk is denoted by $\mathcal{C}_{L,Q,x}^* := \inf_{\alpha \in \mathcal{A}} \mathcal{C}_{L,Q,x}(\alpha)$.

Given a distribution P on $X \times Y$ the inner risks $\mathcal{C}_{L,P(\cdot|x),x}(\alpha)$, $x \in X$, of α can be used to compute the risk $\mathcal{R}_{L,P}(\alpha)$ since (1) immediately gives

$$\mathcal{R}_{L,P}(\alpha) = \int_X \mathcal{C}_{L,P(\cdot|x),x}(\alpha) dP_X(x).$$

Our first goal is to establish the same relation between the minimal inner risks $\mathcal{C}_{L,P(\cdot|x),x}^*$, $x \in X$, and the minimal risk $\mathcal{R}_{L,P}^*$. To this end we define:

Definition 2.4 (Minimizable cost functions) Let $L : X \times Y \times \mathcal{A} \rightarrow [0, \infty]$ be a cost function and P be a distribution on $X \times Y$. We say that L is P -minimizable if for all $\varepsilon > 0$ there exists an $\alpha_\varepsilon \in \mathcal{A}$ such that for all $x \in X$ we have

$$\mathcal{C}_{L,P(\cdot|x),x}(\alpha_\varepsilon) < \mathcal{C}_{L,P(\cdot|x),x}^* + \varepsilon. \quad (2)$$

Note that the above definition in particular ensures $\mathcal{C}_{L,P(\cdot|x),x}^* < \infty$ for all $x \in X$. Now we can establish the following simple but important lemma which computes $\mathcal{R}_{L,P}^*$.

Lemma 2.5 Let P be a distribution on $X \times Y$ and $L : X \times Y \times \mathcal{A} \rightarrow [0, \infty]$ be a P -minimizable cost function. Then $x \mapsto \mathcal{C}_{L,P(\cdot|x),x}^*$ is measurable and we have

$$\mathcal{R}_{L,P}^* = \int_X \mathcal{C}_{L,P(\cdot|x),x}^* dP_X(x).$$

The above lemma shows that the minimal risk $\mathcal{R}_{L,P}^*$ can be achieved by *pointwise minimization* of the inner risks $\mathcal{C}_{L,P(\cdot|x),x}^*$, $x \in X$, which—in general—will be easier to deal with than direct minimization of $\mathcal{R}_{L,P}(\cdot)$. Moreover, for P -minimizable L with $\mathcal{R}_{L,P}^* < \infty$ we have

$$\mathcal{R}_{L,P}(\alpha) - \mathcal{R}_{L,P}^* = \int_X \mathcal{C}_{L,P(\cdot|x),x}(\alpha) - \mathcal{C}_{L,P(\cdot|x),x}^* dP_X(x)$$

for all $\alpha \in \mathcal{A}$. Consequently, we can split the analysis of $\mathcal{R}_{L,P}(\alpha) - \mathcal{R}_{L,P}^*$ into *a)* the analysis of the inner excess risks $\mathcal{C}_{L,P(\cdot|x),x}(\alpha) - \mathcal{C}_{L,P(\cdot|x),x}^*$, $x \in X$, and *b)* the investigation of the integration with respect to P_X . Note, that the benefit of this approach is that the analysis in *a)* only depends on P via the conditional distributions $P(\cdot|x)$ and hence we can consider the excess inner risks $\mathcal{C}_{L,Q,x}(\alpha) - \mathcal{C}_{L,Q,x}^*$ for suitable distributions Q on Y as a template for $\mathcal{C}_{L,P(\cdot|x),x}(\alpha) - \mathcal{C}_{L,P(\cdot|x),x}^*$. This leads to the following definition:

Definition 2.6 Let \mathcal{Q} be a set of distributions on Y . We say that a distribution P on $X \times Y$ is of type \mathcal{Q} if $P(\cdot|x) \in \mathcal{Q}$ for all $x \in X$.

Before we begin with our analysis let us introduce some more notations which will be very useful. The sets containing the elements in \mathcal{A} that “almost” minimize the inner risk in x are denoted by

$$\mathcal{M}_{L,Q,x}(\varepsilon) := \{\alpha \in \mathcal{A} : \mathcal{C}_{L,Q,x}(\alpha) < \mathcal{C}_{L,Q,x}^* + \varepsilon\}, \quad \varepsilon \in [0, \infty].$$

For later use we note that we always have $\mathcal{M}_{L,Q,x}(0) = \emptyset$ and $\mathcal{M}_{L,Q,x}(\varepsilon_1) \subset \mathcal{M}_{L,Q,x}(\varepsilon_2)$ for $0 \leq \varepsilon_1 \leq \varepsilon_2 \leq \infty$. Furthermore we have $\mathcal{M}_{L,Q,x}(\varepsilon) \neq \emptyset$ for some $\varepsilon \in (0, \infty]$ if and only if $\mathcal{C}_{L,Q,x}^* < \infty$. In particular, for P -minimizable cost functions the sets $\mathcal{M}_{L,P(\cdot|x),x}(\varepsilon)$, $\varepsilon > 0$, $x \in X$, are non-empty. Finally, the set

$$\mathcal{M}_{L,Q,x}(0^+) := \bigcap_{\varepsilon > 0} \mathcal{M}_{L,Q,x}(\varepsilon)$$

contains those elements in \mathcal{A} that *exactly* minimize $\mathcal{C}_{L,Q,x}(\cdot)$.

Let us now turn to the main goal of this section, namely comparing the excess risks of different cost functions. To this end we first observe that given two cost functions $L_1 : X \times Y \times \mathcal{A} \rightarrow [0, \infty)$ and $L_2 : X \times Y \times \mathcal{A} \rightarrow [0, \infty)$ with

$$\mathcal{M}_{L_2,P(\cdot|x),x}(0^+) \subset \mathcal{M}_{L_1,P(\cdot|x),x}(0^+), \quad x \in X,$$

Lemma 2.5 shows that every exact minimizer $\alpha_{L_2,P}^* \in \mathcal{A}$ of $\mathcal{R}_{L_2,P}(\cdot)$ is also an exact minimizer of $\mathcal{R}_{L_1,P}(\cdot)$. However, exact minimizers do not necessarily exist, and even if they do exist it is rather unlikely that we will find them by a learning procedure. On the other hand, many learning procedures guarantee to find approximate minimizers with high probability, and thus our overall goal is to establish properties similar to the above observation for approximate minimizers. Let us begin with the following definition.

Definition 2.7 (Calibration) Let \mathcal{Q} be a set of distributions on Y , and $L_1 : X \times Y \times \mathcal{A} \rightarrow [0, \infty]$, $L_2 : X \times Y \times \mathcal{A} \rightarrow [0, \infty]$ be two cost functions. We say that L_2 is L_1 -calibrated with respect to \mathcal{Q} if for all $\varepsilon > 0$, $Q \in \mathcal{Q}$, and $x \in X$ there exists a $\delta > 0$ such that $\mathcal{M}_{L_2,Q,x}(\delta) \subset \mathcal{M}_{L_1,Q,x}(\varepsilon)$, i.e. if for all $\alpha \in \mathcal{A}$ we have

$$\mathcal{C}_{L_2,Q,x}(\alpha) < \mathcal{C}_{L_2,Q,x}^* + \delta \quad \implies \quad \mathcal{C}_{L_1,Q,x}(\alpha) < \mathcal{C}_{L_1,Q,x}^* + \varepsilon. \quad (3)$$

Furthermore, we say that L_2 is L_1 -calibrated with respect to a distribution P on $X \times Y$ if L_2 is L_1 -calibrated with respect the set $\{P(\cdot|x) : x \in X\}$.

The above definition says that L_2 is L_1 -calibrated with respect to P if every element $\alpha \in \mathcal{A}$ minimizing $\mathcal{C}_{L_2,P(\cdot|x),x}(\cdot)$ up to δ minimizes $\mathcal{C}_{L_1,P(\cdot|x),x}(\cdot)$ up to ε . Our first main result now shows that a similar relation then often holds for the risks $\mathcal{R}_{L_1,P}(\cdot)$ and $\mathcal{R}_{L_2,P}(\cdot)$.

Theorem 2.8 Let P be a distribution on $X \times Y$ and $L_1 : X \times Y \times \mathcal{A} \rightarrow [0, \infty]$, $L_2 : X \times Y \times \mathcal{A} \rightarrow [0, \infty]$ be two P -minimizable cost functions such that $\mathcal{R}_{L_1,P}^* < \infty$ and $\mathcal{R}_{L_2,P}^* < \infty$. Furthermore assume that there exist a function $b \in L_1(P_X)$ and measurable functions $\delta(\varepsilon, \cdot) : X \rightarrow (0, \infty)$, $\varepsilon > 0$, with

$$\mathcal{C}_{L_1,P(\cdot|x),x}(\alpha) \leq \mathcal{C}_{L_1,P(\cdot|x),x}^* + b(x)$$

and

$$\mathcal{C}_{L_2,P(\cdot|x),x}(\alpha) < \mathcal{C}_{L_2,P(\cdot|x),x}^* + \delta(\varepsilon, x) \quad \implies \quad \mathcal{C}_{L_1,P(\cdot|x),x}(\alpha) < \mathcal{C}_{L_1,P(\cdot|x),x}^* + \varepsilon \quad (4)$$

for all $x \in X$, $\varepsilon > 0$ and $\alpha \in \mathcal{A}$. Then for all $\varepsilon > 0$ there exists a $\delta > 0$ such that for all $\alpha \in \mathcal{A}$ we have

$$\mathcal{R}_{L_2, P}(\alpha) < \mathcal{R}_{L_2, P}^* + \delta \implies \mathcal{R}_{L_1, P}(\alpha) < \mathcal{R}_{L_1, P}^* + \varepsilon. \quad (5)$$

Note that (5) is equivalent to a positive answer of Question 1, and consequently Theorem 2.8 gives a sufficient condition for a asymptotic relationship between different excess risks in the sense of this question. Moreover, Condition (4) implies that L_2 is L_1 -calibrated with respect to P , and in fact, the only difference between this condition and the calibration with respect to P is the *measurability* of $\delta(\varepsilon, \cdot)$ which is necessary for technical reasons. However, we will see later in Theorem 2.23 that for \mathcal{A} consisting of *functions over X* , this measurability is often automatically satisfied. In addition, Theorem 2.23 shows that in this case the L_1 -calibration of L_2 is also a *necessary* condition for (5), and hence this theorem will give us a complete answer to Question 1.

In order to use Theorem 2.8 the main difficulty is usually to find numbers $\delta(\varepsilon, x) > 0$ with $\mathcal{M}_{L_2, P(\cdot|x), x}(\delta(\varepsilon, x)) \subset \mathcal{M}_{L_1, P(\cdot|x), x}(\varepsilon)$, i.e. to ensure the calibration. The following lemma describes an easy yet optimal solution for this task.

Lemma 2.9 (Calibration function) *Let Q be a distribution on Y , and $L_1 : X \times Y \times \mathcal{A} \rightarrow [0, \infty]$, $L_2 : X \times Y \times \mathcal{A} \rightarrow [0, \infty]$ be two cost functions. Then we define the calibration function $\delta_{\max}(\cdot, Q, x) : [0, \infty] \rightarrow [0, \infty]$ of (L_1, L_2) by*

$$\delta_{\max}(\varepsilon, Q, x) := \begin{cases} \inf_{\alpha \in \mathcal{A} \setminus \mathcal{M}_{L_1, Q, x}(\varepsilon)} \mathcal{C}_{L_2, Q, x}(\alpha) - \mathcal{C}_{L_2, Q, x}^* & \text{if } \mathcal{C}_{L_2, Q, x}^* < \infty, \\ \infty & \text{if } \mathcal{C}_{L_2, Q, x}^* = \infty \end{cases} \quad (6)$$

for all $\varepsilon \in [0, \infty]$. Then for all $\varepsilon \in [0, \infty]$ we have:

- i) $\mathcal{M}_{L_2, Q, x}(\delta_{\max}(\varepsilon, Q, x)) \subset \mathcal{M}_{L_1, Q, x}(\varepsilon)$.
- ii) $\mathcal{M}_{L_2, Q, x}(\delta) \not\subset \mathcal{M}_{L_1, Q, x}(\varepsilon)$ whenever $\delta > \delta_{\max}(\varepsilon, Q, x)$.

In addition, if both $\mathcal{C}_{L_1, Q, x}^* < \infty$ and $\mathcal{C}_{L_2, Q, x}^* < \infty$, then for all $\alpha \in \mathcal{A}$ we have

$$\delta_{\max}(\mathcal{C}_{L_1, Q, x}(\alpha) - \mathcal{C}_{L_1, Q, x}^*, Q, x) \leq \mathcal{C}_{L_2, Q, x}(\alpha) - \mathcal{C}_{L_2, Q, x}^*. \quad (7)$$

Note that in some cases we have to distinguish between different calibration functions and hence we occasionally use the notation $\delta_{\max, L_1, L_2}(\cdot, Q, x) := \delta_{\max}(\cdot, Q, x)$ for the calibration function $\delta_{\max}(\cdot, Q, x)$ of (L_1, L_2) .

It turns out that in most practical situations the computation of the calibration function is a straightforward exercise. In particular, if the surrogate loss is *convex*¹ in the sense of the following definition then this computation is rather easy as we will see below:

Definition 2.10 *A cost function $L : X \times Y \times \mathbb{R} \rightarrow [0, \infty]$ is called (strictly) convex if $L(x, y, \cdot) : \mathbb{R} \rightarrow [0, \infty]$ is (strictly) convex for all $x \in X$ and $y \in Y$.*

As already indicated, the calibration function for convex surrogates can be easily computed. This is stated in the following result:

Lemma 2.11 *Let Q be a distribution on Y , $x \in X$, $\varepsilon > 0$, and $L_1 : X \times Y \times \mathbb{R} \rightarrow [0, \infty]$ be a cost function such that $\mathcal{M}_{L_1, Q, x}(\varepsilon)$ is an interval. Moreover, let $L_2 : X \times Y \times \mathbb{R} \rightarrow [0, \infty]$ be a convex*

¹Basic definitions and properties regarding convexity can be found in the appendix.

cost function such that $\mathcal{C}_{L_2, Q, x}(\cdot) : \mathbb{R} \rightarrow [0, \infty)$ is continuous. If $\mathcal{M}_{L_2, Q, x}(0^+) \subset \mathcal{M}_{L_1, Q, x}(0^+)$ then we have

$$\delta_{\max}(\varepsilon, Q, x) = \min \left\{ \mathcal{C}_{L_2, Q, x}(\sup \mathcal{M}_{L_1, Q, x}(\varepsilon)), \mathcal{C}_{L_2, Q, x}(\inf \mathcal{M}_{L_1, Q, x}(\varepsilon)) \right\} - \mathcal{C}_{L_2, Q, x}^*, \quad (8)$$

where $\mathcal{C}_{L_2, Q, x}(\pm\infty) := \infty$.

Let us now investigate Question 2. To this end observe that in some sense inequalities between the involved excess risks are readily available by Theorem 2.8: indeed, for $\varepsilon > 0$ let $\delta(\varepsilon) > 0$ be a real number such that implication (5) holds for all $\alpha \in \mathcal{A}$. Furthermore, we define $\delta(0) := 0$. For $\alpha \in \mathcal{A}$ with $\varepsilon := \mathcal{R}_{L_1, P}(\alpha) - \mathcal{R}_{L_1, P}^* > 0$ we then have $\mathcal{R}_{L_2, P}(\alpha) - \mathcal{R}_{L_2, P}^* \geq \delta(\varepsilon)$, or in other words

$$\delta(\mathcal{R}_{L_1, P}(\alpha) - \mathcal{R}_{L_1, P}^*) \leq \mathcal{R}_{L_2, P}(\alpha) - \mathcal{R}_{L_2, P}^*. \quad (9)$$

Furthermore, for $\alpha \in \mathcal{A}$ with $\mathcal{R}_{L_1, P}(\alpha) - \mathcal{R}_{L_1, P}^* = 0$ this inequality is satisfied by our definition $\delta(0) = 0$, and consequently (9) is satisfied for all $\alpha \in \mathcal{A}$. Inverting this inequality then gives a positive answer to Question 2. However, the proof of Theorem 2.8 does unfortunately not provide a constructive way to find a value for $\delta(\varepsilon)$, and hence our next aim is to describe conditions on L_1 , L_2 , and P which will allow us to easily establish inequalities in many situations. To this end we begin with the following definition.

Definition 2.12 (Fenchel-Legendre biconjugate) Let $I \subset \mathbb{R}$ be either $I := [0, B]$ for some $B > 0$, or $I := [0, \infty)$. Furthermore, let $\delta : I \rightarrow [0, \infty]$ be a function. Then the Fenchel-Legendre bi-conjugate $\delta^{**} : I \rightarrow [0, \infty]$ of δ is the function defined by

$$\text{Epi } \delta^{**} = \overline{\text{co Epi } \delta},$$

where $\text{Epi } \delta := \{(\varepsilon, y) \in I \times [0, \infty] : \delta(\varepsilon) \leq y\}$ denotes the epigraph of δ and $\text{co } A$ denotes the convex hull of a set A .

Note that the Fenchel-Legendre bi-conjugate δ^{**} which is a well-known tool in convex analysis (see e.g. [23]), is the largest convex function f satisfying $f \leq \delta$. Moreover, for technical reasons we use the convention $\delta^{**}(\infty) := \lim_{\varepsilon \rightarrow \infty} \delta^{**}(\varepsilon)$ for functions $\delta : [0, \infty) \rightarrow [0, \infty)$.

Now we are well prepared to establish our first inequalities between excess risks of different cost functions.

Theorem 2.13 Let P be a distribution on $X \times Y$ and $L_1 : X \times Y \times \mathcal{A} \rightarrow [0, \infty]$, $L_2 : X \times Y \times \mathcal{A} \rightarrow [0, \infty]$ be two P -minimizable cost functions with $\mathcal{R}_{L_1, P}^* < \infty$ and $\mathcal{R}_{L_2, P}^* < \infty$. We write

$$B_\alpha := \text{ess-sup}_{x \in X} \mathcal{C}_{L_1, P(\cdot|x), x}(\alpha) - \mathcal{C}_{L_1, P(\cdot|x), x}^*$$

for all $\alpha \in \mathcal{A}$. Furthermore, let $\delta : [0, \infty] \rightarrow [0, \infty]$ be a function with $\delta(0) = 0$ and

$$\mathcal{M}_{L_2, P(\cdot|x), x}(\delta(\varepsilon)) \subset \mathcal{M}_{L_1, P(\cdot|x), x}(\varepsilon) \quad (10)$$

for all $x \in X$ and all $\varepsilon \in [0, \infty]$. Then for all $\alpha \in \mathcal{A}$ we have

$$\delta_{B_\alpha}^{**}(\mathcal{R}_{L_1, P}(\alpha) - \mathcal{R}_{L_1, P}^*) \leq \mathcal{R}_{L_2, P}(\alpha) - \mathcal{R}_{L_2, P}^*, \quad (11)$$

where $\delta_{B_\alpha}^{**} : [0, B_\alpha] \rightarrow [0, \infty]$ denotes the Fenchel-Legendre biconjugate of the restriction $\delta|_{[0, B_\alpha]}$.

Remark 2.14 If the function δ in the above theorem is increasing with $\delta(\varepsilon) > 0$ for all $\varepsilon > 0$ and we have $B_\alpha < \infty$ then Lemma A.6 shows that its bi-conjugate satisfies $\delta_{B_\alpha}^{**}(\varepsilon) > 0$ for all $\varepsilon \in (0, B_\alpha]$. Moreover, the bi-conjugate is always convex, and since $\delta_{B_\alpha}^{**}(0) = 0$ it is also *strictly* increasing and thus injective. Consequently, if $\delta_{B_\alpha}^{**}$ is finite, then it has a continuous inverse, and hence we have found a positive answer to Question 2. However, in general this is no longer true if $B_\alpha = \infty$, as e.g. the case $\delta(\varepsilon) = \sqrt{\varepsilon}$, $\varepsilon \geq 0$, shows (see also [33, Prop. A.5] for conditions guaranteeing $\delta^{**} > 0$).

We will see in Theorem 2.18 that Condition (10) is also *necessary* if one wants to have an inequality between the two excess risks that is *independent* of the specific distribution P . Of course, from a machine learning perspective this independence is a highly desired property since in general the data-generating distribution P is assumed to be unknown. Finally, we will see in Section 3 that for binary classification the inequalities established in Theorem 2.13 coincide with the inequalities found in [1]. Since their inequalities are sharp if one does not make additional assumptions on P , we see that the inequalities of Theorem 2.13 cannot be improved in general. However, note that under certain conditions on P there are substantial improvements possible, and we will present examples of such improvements later when dealing with more specific cost functions.

If the function δ in the above theorem satisfies $\delta(\varepsilon) > 0$ for all $\varepsilon > 0$ then condition (10) is a *uniform* version of the notion of calibration. Let us describe this situation in the following definition:

Definition 2.15 (Uniform calibration) *Let \mathcal{Q} be a set of distributions on Y , and $L_1 : X \times Y \times \mathcal{A} \rightarrow [0, \infty)$, $L_2 : X \times Y \times \mathcal{A} \rightarrow [0, \infty)$ be two cost functions. We say that L_2 is uniformly L_1 -calibrated with respect to \mathcal{Q} if for all $\varepsilon > 0$ there exists a $\delta > 0$ such that $\mathcal{M}_{L_2, Q, x}(\delta) \subset \mathcal{M}_{L_1, Q, x}(\varepsilon)$ for all $Q \in \mathcal{Q}$ and $x \in X$, i.e. if for all $Q \in \mathcal{Q}$, $x \in X$, and $\alpha \in \mathcal{A}$ we have*

$$\mathcal{C}_{L_2, Q, x}(\alpha) < \mathcal{C}_{L_2, Q, x}^* + \delta \quad \implies \quad \mathcal{C}_{L_1, Q, x}(\alpha) < \mathcal{C}_{L_1, Q, x}^* + \varepsilon.$$

Furthermore, we say that L_2 is uniformly L_1 -calibrated with respect to a distribution P on $X \times Y$ if L_2 is uniformly L_1 -calibrated with respect to the set $\{P(\cdot|x) : x \in X\}$.

Obviously the above definition guarantees that L_1 is P -minimizable whenever L_2 is so. Now observe that the larger the function δ satisfying (10) is, the better inequality (11) becomes. Given a set \mathcal{Q} of distributions on Y let us consequently consider the best possible candidate function, that is

$$\delta_{\max}(\varepsilon, \mathcal{Q}) := \sup \left\{ \delta \geq 0 : \mathcal{M}_{L_2, Q, x}(\delta) \subset \mathcal{M}_{L_1, Q, x}(\varepsilon) \text{ for all } x \in X, Q \in \mathcal{Q} \right\}, \quad \varepsilon \in [0, \infty]. \quad (12)$$

Obviously, L_2 is uniformly L_1 -calibrated with respect to \mathcal{Q} if and only if $\delta_{\max}(\varepsilon, \mathcal{Q}) > 0$ for all $\varepsilon > 0$. Moreover, given a distribution P on $X \times Y$ we write $\delta_{\max}(\varepsilon, P) := \delta_{\max}(\varepsilon, \{P(\cdot|x) : x \in X\})$. Assuming that P is of type \mathcal{Q} we then have $\delta_{\max}(\varepsilon, \mathcal{Q}) \leq \delta_{\max}(\varepsilon, P)$ and consequently, the next lemma shows that the distribution-independent function $\delta_{\max}(\cdot, \mathcal{Q})$ can be used in Theorem 2.13.²

Lemma 2.16 *Let \mathcal{Q} be a set of distributions on Y , and $L_1 : X \times Y \times \mathcal{A} \rightarrow [0, \infty)$, $L_2 : X \times Y \times \mathcal{A} \rightarrow [0, \infty)$ be two cost functions. Then for all $x \in X$ and all $\varepsilon > 0$ we have*

$$\mathcal{M}_{L_2, Q, x}(\delta_{\max}(\varepsilon, \mathcal{Q})) \subset \mathcal{M}_{L_1, Q, x}(\varepsilon).$$

In order to use $\delta_{\max}(\cdot, P)$ or $\delta_{\max}(\cdot, \mathcal{Q})$ in concrete examples we also need a simple method to calculate these quantities. This is done in the following lemma:

²In Theorem 2.26 we will see that rather natural conditions, using $\delta_{\max}(\cdot, \mathcal{Q})$ does not give worse inequalities than using $\delta_{\max}(\varepsilon, P)$.

Lemma 2.17 *Let \mathcal{Q} be a set of distributions on Y , and $L_1 : X \times Y \times \mathcal{A} \rightarrow [0, \infty)$, $L_2 : X \times Y \times \mathcal{A} \rightarrow [0, \infty)$ be two cost functions. Then for all $\varepsilon > 0$ we have*

$$\delta_{\max}(\varepsilon, \mathcal{Q}) = \inf_{\substack{Q \in \mathcal{Q} \\ x \in X}} \delta_{\max}(\varepsilon, Q, x) .$$

Theorem 2.13 states that for distributions P of type \mathcal{Q} we can find inequalities between the excess risks of uniformly calibrated cost functions. The following theorem shows that the uniform calibration is also *necessary* for such inequalities.

Theorem 2.18 *Let $L_1 : X \times Y \times \mathcal{A} \rightarrow [0, \infty)$ and $L_2 : X \times Y \times \mathcal{A} \rightarrow [0, \infty)$ be cost functions, and \mathcal{Q} be class of distributions on Y with $\mathcal{C}_{L_1, Q, x}^* < \infty$ and $\mathcal{C}_{L_2, Q, x}^* < \infty$ for all $x \in X$, $Q \in \mathcal{Q}$. Furthermore, let $\delta : [0, \infty] \rightarrow [0, \infty]$ be an increasing function with $\delta(0) = 0$ and $\delta(\varepsilon) > 0$ for all $\varepsilon > 0$. If for all distributions P of type \mathcal{Q} satisfying $\mathcal{R}_{L_1, P}^* < \infty$ and $\mathcal{R}_{L_2, P}^* < \infty$, and all measurable $\alpha \in \mathcal{A}$ we have*

$$\delta(\mathcal{R}_{L_1, P}(\alpha) - \mathcal{R}_{L_1, P}^*) \leq \mathcal{R}_{L_2, P}(\alpha) - \mathcal{R}_{L_2, P}^* ,$$

then L_2 is uniformly L_1 -calibrated with respect to \mathcal{Q} .

It will turn out that in many situations such as regression we also need inequalities between the excess risks of calibrated cost functions which are not uniformly calibrated. The following theorem presents a method for dealing with this situation.

Theorem 2.19 *Let P be a distribution on $X \times Y$ and $L_1 : X \times Y \times \mathcal{A} \rightarrow [0, \infty)$, $L_2 : X \times Y \times \mathcal{A} \rightarrow [0, \infty)$ be P -minimizable cost functions. Assume that there exist $p \in (0, \infty]$, $q \geq \frac{p+1}{p}$, and a $b : X \rightarrow [0, \infty]$ such that*

$$\delta_{\max}(\varepsilon, P(\cdot|x), x) \geq \varepsilon^q b(x) , \quad \varepsilon > 0, x \in X \quad (13)$$

and $b^{-1} \in L_p(P_X)$. Then for all $\alpha \in \mathcal{A}$ we have

$$\mathcal{R}_{L_1, P}(\alpha) - \mathcal{R}_{L_1, P}^* \leq \|b^{-1}\|_{L_p(P_X)}^{\frac{1}{q}} \left(\mathcal{R}_{L_2, P}(\alpha) - \mathcal{R}_{L_2, P}^* \right)^{\frac{1}{q}} .$$

Remark 2.20 The above condition $b^{-1} \in L_p(P_X)$ measures how much the calibration function $\delta_{\max}(\varepsilon, P(\cdot|x), x)$ violates a uniform lower bound of the form $\delta_{\max}(\varepsilon, P(\cdot|x), x) \geq c\varepsilon^q$, $\varepsilon \in [0, \infty]$. Indeed, the larger we can choose p in Condition (13) the more the shape of b is away from the critical level 0, and thus the closer Condition (13) is to a uniform lower bound. In the extremal case $p = \infty$, Condition (13) actually becomes a uniform bound, and the inequality of Theorem 2.19 equals the inequality of Theorem 2.13.

2.2 Loss Functions

In the previous subsection we have not specified the hypothesis space \mathcal{A} . In this subsection we will now investigate the important case where this space consists of measurable functions. We begin with the following fundamental definition:

Definition 2.21 (Loss function) *Let \mathcal{T} be a Polish space. Then a function $L : X \times Y \times \mathcal{T} \rightarrow [0, \infty]$ is called a loss function if it is measurable.*

Obviously, every loss function L is a cost function, but in general the converse is not true. Moreover, a loss function L also induces a cost function acting on the space $\mathcal{M}(X, \mathcal{T})$ of all measurable functions $f : X \rightarrow \mathcal{T}$. Indeed, for $\mathcal{A} := \mathcal{M}(X, \mathcal{T})$ the mapping

$$\begin{aligned} \hat{L} : X \times Y \times \mathcal{A} &\rightarrow [0, \infty) \\ (x, y, f) &\mapsto L(x, y, f(x)) \end{aligned}$$

is a cost function since $(x, y) \mapsto L(x, y, f(x))$ is measurable for all $f \in \mathcal{A}$. Note that the inner risks of L and \hat{L} are related by

$$\mathcal{C}_{\hat{L}, Q, x}(f) = \int_Y L(x, y, f(x)) dQ(y) = \mathcal{C}_{L, Q, x}(f(x)), \quad (14)$$

and consequently we have $\mathcal{C}_{\hat{L}, Q, x}^* = \mathcal{C}_{L, Q, x}^*$. In the following we are *only* interested in the excess risks of the induced cost function \hat{L} , and therefore we write in a slight abuse of notations

$$\mathcal{R}_{L, P}(f) := \mathcal{R}_{\hat{L}, P}(f)$$

for measurable functions $f : X \rightarrow \mathcal{T}$. Analogously, we define $\mathcal{R}_{L, P}^* := \mathcal{R}_{\hat{L}, P}^*$. Now recall that all the major results of the previous subsection required that the involved cost functions are P -minimizable. Having a loss function L we consequently need to ensure that its induced cost function \hat{L} is P -minimizable. This is done in the following lemma whose main difficulty is to ensure the measurability statements.

Theorem 2.22 *Let $L : X \times Y \times \mathcal{T} \rightarrow [0, \infty]$ be a loss function and P be a distribution on $X \times Y$. Then the following statements are true:*

- i) \hat{L} is P -minimizable if and only if $\mathcal{C}_{L, P(\cdot|x), x}^* < \infty$ and all $x \in \mathcal{X}$.
- ii) If we have $\mathcal{M}_{L, P(\cdot|x), x}(0^+) \neq \emptyset$ for all $x \in \mathcal{X}$ then there exists a measurable function $f_{L, P}^* : X \rightarrow \mathcal{T}$ with $\mathcal{C}_{L, P(\cdot|x), x}(f_{L, P}^*(x)) = \mathcal{C}_{L, P(\cdot|x), x}^*$ for all $x \in X$, and consequently we have

$$\mathcal{R}_{L, P}(f_{L, P}^*) = \mathcal{R}_{L, P}^*.$$

Now recall the the key quantity of the previous section was the calibration function. Since we are only interested in the excess risks of the cost functions \hat{L}_1 and \hat{L}_2 induced by the loss functions L_1 and L_2 we need to calculate the calibration function of (\hat{L}_1, \hat{L}_2) . To this end assume first that we have $\mathcal{C}_{L_2, Q, x}^* < \infty$. Then Equation (14) gives

$$\delta_{\max, L_1, L_2}(\varepsilon, Q, x) = \inf_{\substack{t \in \mathcal{T} \\ t \notin \mathcal{M}_{L_1, Q, x}(\varepsilon)}} \mathcal{C}_{L_2, Q, x}(t) - \mathcal{C}_{L_2, Q, x}^* = \delta_{\max, \hat{L}_1, \hat{L}_2}(\varepsilon, Q, x),$$

where in the last equality we used that $f \equiv t$ is a measurable function with $f(x) = t$. Moreover, if $\mathcal{C}_{L_2, Q, x}^* = \infty$ the above observation is also true, and consequently it suffices to investigate the calibration function $\delta_{\max, L_1, L_2}(\varepsilon, Q, x)$ when dealing with the pair (\hat{L}_1, \hat{L}_2) .

We have already seen in Lemma 2.9 that the calibration function satisfies condition (4), however, in order to use the calibration function in Theorem 2.8 we also need to know that it is measurable. The following theorem completely resolves this issue, and in addition it *characterizes* the distributions for which implication (5) hold:

Theorem 2.23 *Let $L_1 : X \times Y \times \mathcal{T} \rightarrow [0, \infty]$ and $L_2 : X \times Y \times \mathcal{T} \rightarrow [0, \infty]$ be loss functions, and P be a distribution on $X \times Y$ such that $\mathcal{R}_{L_1, P}^* < \infty$ and $\mathcal{R}_{L_2, P}^* < \infty$. Then*

$$x \mapsto \delta_{\max}(\varepsilon, P(\cdot|x), x)$$

is measurable for all $\varepsilon > 0$. Furthermore, if for all $\varepsilon > 0$ there exists a $\delta > 0$ with

$$\mathcal{R}_{L_2, P}(f) < \mathcal{R}_{L_2, P}^* + \delta \quad \implies \quad \mathcal{R}_{L_1, P}(f) < \mathcal{R}_{L_1, P}^* + \varepsilon$$

for all measurable $f : X \rightarrow \mathcal{T}$, then we have $P_X(\{x \in X : \delta_{\max}(\varepsilon, x) = 0\}) = 0$ for all $\varepsilon > 0$.

The above theorem shows that the L_1 -calibration is *necessary* for L_2 being a reasonable surrogate loss in the sense of Question 1. Consequently, the calibration function will be our major tool when investigating specific learning problems in Section 3.

2.2.1 Regular Loss Functions

General loss functions can explicitly depend on the input variable x , and hence so do the derived quantities like the inner risks. However, many important loss functions are actually independent of x , and since the theory becomes substantially simpler for such losses we now briefly consider them. Let us begin with the following definition:

Definition 2.24 (Regular loss functions) *Let \mathcal{T} be a Polish space. A function $L : Y \times \mathcal{T} \rightarrow [0, \infty]$ is called a regular loss function if it is measurable.*

Formally, regular loss functions are not loss functions, however, it is obvious that every regular loss function L induces a loss function \bar{L} via $\bar{L}(x, y, t) := L(y, t)$. In the following we *always* identify L with \bar{L} . Now note that the inner risk of \bar{L} is

$$\mathcal{C}_{\bar{L}, Q, x}(t) = \int_Y L(y, t) dQ(y),$$

i.e. it is *independent of x* . Moreover, the same is obviously true for the derived quantities $\mathcal{C}_{L, Q, x}^*$, $\mathcal{M}_{L, Q, x}(\varepsilon)$, and therefore we usually use the shorter notations

$$\mathcal{C}_{L, Q}(t) := \mathcal{C}_{\bar{L}, Q, x}(t), \quad \mathcal{C}_{L, Q}^* := \mathcal{C}_{L, Q, x}^*, \quad \text{and} \quad \mathcal{M}_{L, Q}(\varepsilon) := \mathcal{M}_{L, Q, x}(\varepsilon),$$

where $x \in X$ is an arbitrary element. Furthermore, note that if we have two regular loss functions then the corresponding calibration function $\delta_{\max}(\varepsilon, Q, x)$ is independent of x , too, and hence we analogously write $\delta_{\max}(\varepsilon, Q) := \delta_{\max}(\varepsilon, Q, x)$ for some $x \in X$. Finally note that the often imposed condition $\mathcal{C}_{L, Q, x}^* < \infty$ is also independent of x . This justifies the following definition:

Definition 2.25 *Let \mathcal{Q} be a class of distributions on Y and $L : Y \times \mathcal{T} \rightarrow [0, \infty)$ be a regular loss function. Then we write*

$$\mathcal{Q}(L) := \{Q \in \mathcal{Q} : \mathcal{C}_{L, Q}^* < \infty\}.$$

Obviously, all theorems we have formulated for cost or loss functions also hold for regular loss functions. In particular, for uniformly calibrated regular loss functions Theorem 2.13 can be used to establish inequalities between the corresponding excess risks. Now note that for distributions P of type \mathcal{Q} we always have $\delta_{\max}(\varepsilon, P) \leq \delta_{\max}(\varepsilon, \mathcal{Q})$, and consequently one may think that using $\delta_{\max}(\varepsilon, P)$ instead of $\delta_{\max}(\varepsilon, \mathcal{Q})$ leads to sharper inequalities in Theorem 2.13. The following result shows that this intuition is usually false if the set \mathcal{Q} is suitably chosen.

Theorem 2.26 *Let $L_1 : Y \times \mathcal{T} \rightarrow [0, \infty]$, $L_2 : Y \times \mathcal{T} \rightarrow [0, \infty]$ be regular loss functions and \mathcal{Q} be a set of distributions on Y with $\mathcal{Q} = \mathcal{Q}(L_2) = \mathcal{Q}(L_1)$. Furthermore assume that there is a distribution μ on X for which there exist mutually disjoint subsets $A_n \subset X$, $n \in \mathbb{N}$, with $\mu(A_n) > 0$ for all $n \in \mathbb{N}$. Then there exists a distribution P of type \mathcal{Q} with $P_X = \mu$ such that for all $\varepsilon > 0$ we have*

$$\delta_{\max}(\varepsilon, P) = \delta_{\max}(\varepsilon, \mathcal{Q}). \tag{15}$$

2.2.2 Unsupervised Loss Functions

In the previous part of the work we considered loss functions whose inner risks are independent of x . Formally we can also consider loss functions whose inner risks are independent of Q . Before we motivate such losses let us first make a precise definition:

Definition 2.27 (Unsupervised loss functions) *Let \mathcal{T} be a Polish space. A function $L : X \times \mathcal{T} \rightarrow [0, \infty]$ is called an unsupervised loss function if it is measurable.*

Like regular loss functions, unsupervised loss functions are not loss functions. However, it is obvious that every unsupervised loss function L induces a loss function \bar{L} via $\bar{L}(x, y, t) := L(x, t)$. In the following we *always* identify L with \bar{L} . Now note that the inner risk of \bar{L} is

$$\mathcal{C}_{\bar{L}, Q, x}(t) = L(x, t),$$

i.e. it is *independent of Q* . Moreover, the quantities $\mathcal{C}_{L, Q, x}^*$, $\mathcal{M}_{L, Q, x}(\varepsilon)$ obviously share this property. In the following we hence use the shorthands

$$\mathcal{C}_{L, x}(t) := \mathcal{C}_{L, Q, x}(t), \quad \mathcal{C}_{L, x}^* := \mathcal{C}_{L, Q, x}^*, \quad \text{and} \quad \mathcal{M}_{L, x}(\varepsilon) := \mathcal{M}_{L, Q, x}(\varepsilon),$$

where Q is an arbitrary distribution. Note that these definitions in particular give

$$\mathcal{C}_{L, x}^* = \inf_{t \in \mathcal{T}} L(x, t). \tag{16}$$

At a first glance it seems rather odd to consider unsupervised loss functions since their associated (inner) risks $\mathcal{C}_{L, P(\cdot|x), x}(\cdot)$ do not depend on the conditional probabilities $P(\cdot|x)$, and hence they apparently do not define a reasonable learning goal. However, as we will see later, the key idea of unsupervised loss functions is to encode the dependence of $P(\cdot|x)$ on the first input variable x of L , so that we actually have the possibility to encode dependencies that are *not* of the form of an inner risk. This fact will be particularly important when considering unsupervised learning goals such as density level detection and density estimation. In addition unsupervised loss functions will also serve as a technical tool when considering general questions on loss functions in e.g. Subsection 2.3. In this regard the following class of unsupervised loss functions is of particular interest:

Definition 2.28 (Detection loss functions) *Let $A \subset X \times \mathcal{T}$ be a measurable subset and $h : X \rightarrow [0, \infty)$ be a measurable function. Then we call $L : X \times \mathcal{T} \rightarrow [0, \infty)$ a detection loss function with respect to (A, h) if*

$$L(x, t) = \mathbf{1}_A(x, t) h(x), \quad x \in X, t \in \mathcal{T}.$$

Every detection loss function is obviously measurable and hence an unsupervised loss function. In addition, for $x \in X$ and $t \in \mathcal{T}$ our above notations immediately show

$$\mathcal{C}_{L, x}(t) - \mathcal{C}_{L, x}^* = \begin{cases} 0 & \text{if } A(x) := \{t' \in \mathcal{T} : (x, t') \in A\} = \mathcal{T} \\ \mathbf{1}_A(x, t) h(x) & \text{else.} \end{cases} \tag{17}$$

Since detection loss functions will play an important role for *both* supervised and unsupervised learning scenarios let us now establish some specific results for them. We begin with the following simple theorem, which directly follows from Theorem 2.8, Theorem 2.22, and Theorem 2.23:

Theorem 2.29 *Let $L_1 : X \times \mathcal{T} \rightarrow [0, \infty)$ be a detection loss function with respect to some (A, h) , $L_2 : X \times Y \times \mathcal{T} \rightarrow [0, \infty]$ be a loss function, and \mathcal{Q} be a class of distributions on Y . Then the following statements are equivalent:*

i) L_2 is L_1 -calibrated with respect to \mathcal{Q} .

ii) For all distributions P of type \mathcal{Q} with $h \in L_1(P_X)$ and $\mathcal{R}_{L_2,P}^* < \infty$, and all $\varepsilon \in (0, \infty]$ there exists a $\delta \in (0, \infty]$ such that for all measurable $f : X \rightarrow \mathbb{R}$ we have

$$\mathcal{R}_{L_2,P}(f) < \mathcal{R}_{L_2,P}^* + \delta \implies \mathcal{R}_{L_1,P}(f) < \mathcal{R}_{L_1,P}^* + \varepsilon.$$

If L_2 is actually uniformly L_1 -calibrated with respect to P , then it is straightforward to establish inequalities by Theorem 2.13. However, using the specific form of detection losses one can often improve these inequalities as we will discuss after the following rather general theorem:

Theorem 2.30 Let $L_1 : X \times \mathcal{T} \rightarrow [0, \infty)$ be a detection loss function with respect to (A, h) , $L_2 : X \times Y \times \mathcal{T} \rightarrow [0, \infty]$ be a loss function, and P be a distribution on $X \times Y$ with $\mathcal{R}_{L_1,P}^* < \infty$, $\mathcal{R}_{L_2,P}^* < \infty$. For $s > 0$ we write

$$B(s) := \left\{ x \in X : A(x) \neq \mathcal{T} \text{ and } \delta_{\max}(h(x), P(\cdot|x), x) < s h(x) \right\}.$$

If there are constants $c > 0$ and $\alpha \in (0, \infty]$ with

$$\int \mathbf{1}_{B(s)} h dP_X \leq (cs)^\alpha \quad (18)$$

for all $s > 0$, then for all measurable $f : X \rightarrow \mathcal{T}$ we have

$$\mathcal{R}_{L_1,P}(f) - \mathcal{R}_{L_1,P}^* \leq 2c^{\frac{\alpha}{\alpha+1}} (\mathcal{R}_{L_2,P}(f) - \mathcal{R}_{L_2,P}^*)^{\frac{\alpha}{\alpha+1}}.$$

The above theorem improves the inequalities we obtained for general target losses L_1 in various cases. The following two remarks illustrate this:

Remark 2.31 Let L_1 be a detection loss with $h = \mathbf{1}_X$, and assume that Condition (13) is satisfied for some $b : X \rightarrow [0, \infty]$ with $b^{-1} \in L_p(P_X)$ and $q \geq \frac{p+1}{p}$. Then Theorem 2.19 gives

$$\mathcal{R}_{L_1,P}(f) - \mathcal{R}_{L_1,P}^* \leq \|b^{-1}\|_{L_p(P_X)}^{\frac{1}{q}} (\mathcal{R}_{L_2,P}(f) - \mathcal{R}_{L_2,P}^*)^{\frac{1}{q}}. \quad (19)$$

However, we also have $B(s) \subset \{x \in X : b(x) < s\}$, and since $b^{-1} \in L_p(P_X)$ implies $P_X(\{x \in X : b(x) < s\}) \leq (\|b^{-1}\|_p s)^p$, we find (18) for $c := \|b^{-1}\|_{L_p(P_X)}$ and $\alpha := p$. Consequently, Theorem 2.30 yields

$$\mathcal{R}_{L_1,P}(f) - \mathcal{R}_{L_1,P}^* \leq 2\|b^{-1}\|_{L_p(P_X)}^{\frac{p}{p+1}} (\mathcal{R}_{L_2,P}(f) - \mathcal{R}_{L_2,P}^*)^{\frac{p}{p+1}}. \quad (20)$$

Now recall that Theorem 2.19 required $\frac{1}{q} \leq \frac{p}{p+1}$, and therefore (20) is sharper than (19) whenever the excess risk $\mathcal{R}_{L_2,P}(f) - \mathcal{R}_{L_2,P}^*$ is sufficiently small.

Remark 2.32 Let L_1 be a detection loss and L_2 be a loss that is uniformly L_1 -calibrated with respect to some class \mathcal{Q} of distributions. If $\delta_{\max}(\cdot, \mathcal{Q}) \geq \varepsilon^q$ for some $q > 1$ and all $\varepsilon \geq 0$ then Theorem 2.13 gives

$$\mathcal{R}_{L_1,P}(f) - \mathcal{R}_{L_1,P}^* \leq (\mathcal{R}_{L_2,P}(f) - \mathcal{R}_{L_2,P}^*)^{\frac{1}{q}} \quad (21)$$

for all measurable $f : X \rightarrow \mathbb{R}$. However, it holds $B(s) \subset \{x \in X : 0 < h(x) < s^{1/(q-1)}\}$, and consequently, if we have

$$P_X(\{x \in X : 0 < h(x) < s\}) \leq (Cs)^\beta \quad (22)$$

for some constants $C > 0$, $\beta \in (0, \infty]$, and all $s > 0$, then it is easy to check that (18) is satisfied for $c = C^{\frac{\beta q - \beta}{\beta + 1}}$ and $\alpha := \frac{\beta + 1}{q - 1}$. Theorem 2.30 then yields

$$\mathcal{R}_{L_1,P}(f) - \mathcal{R}_{L_1,P}^* \leq 2C^{\frac{\beta q - \beta}{\beta + 1}} (\mathcal{R}_{L_2,P}(f) - \mathcal{R}_{L_2,P}^*)^{\frac{\beta + 1}{\beta + q}}. \quad (23)$$

Now note that we always have $\frac{\beta + 1}{\beta + q} \geq \frac{1}{q}$, and thus (23) is sharper than (21) whenever the excess risk $\mathcal{R}_{L_2,P}(f) - \mathcal{R}_{L_2,P}^*$ is sufficiently small.

2.3 Template loss functions and self-calibrated loss functions

Given a loss function L and a distribution P such that an *exact* minimizer $f_{L,P}^*$ of $\mathcal{R}_{L,P}(\cdot)$ exists one may ask whether, and in which sense, ε -approximate minimizers f_ε of $\mathcal{R}_{L,P}(\cdot)$ approximate $f_{L,P}^*$. The goal of this subsection is to provide some general answers to these questions.

Let us begin with the following definition which introduces another important class of loss functions:

Definition 2.33 (Template loss) *Let \mathcal{Q} be a class of distributions on Y . Then we call a function $L : \mathcal{Q} \times \mathbb{R} \rightarrow [0, \infty]$ a template loss function if for all \mathcal{Q} -type distributions P on $X \times Y$ the P -instance L_P of L defined by*

$$\begin{aligned} L_P : X \times \mathbb{R} &\rightarrow [0, \infty] \\ (x, t) &\mapsto L(P(\cdot|x), t), \end{aligned} \tag{24}$$

is measurable.

In the following template losses will often serve as a powerful tool since they combine properties of both regular and unsupervised loss functions. To make this precise let $L : \mathcal{Q} \times \mathbb{R} \rightarrow [0, \infty]$ be a template loss. We define

$$\mathcal{C}_{L,Q}(t) := L(Q, t), \quad \text{and} \quad \mathcal{C}_{L,Q}^* := \inf_{t' \in \mathbb{R}} L(Q, t')$$

for $Q \in \mathcal{Q}$, $t \in \mathbb{R}$, and write $\mathcal{M}_{L,Q}(\varepsilon) := \{t \in \mathbb{R} : \mathcal{C}_{L,Q}(t) < \mathcal{C}_{L,Q}^* + \varepsilon\}$, $Q \in \mathcal{Q}$, $\varepsilon > 0$. Moreover, given a regular surrogate loss L_2 we define the calibration function $\delta_{\max}(\cdot, Q) : [0, \infty] \rightarrow [0, \infty]$ by

$$\delta_{\max}(\varepsilon, Q) := \delta_{\max, L, L_2}(\varepsilon, Q) := \inf_{\substack{t \in T \\ t \notin \mathcal{M}_{L,Q}(\varepsilon)}} \mathcal{C}_{L_2,Q}(t) - \mathcal{C}_{L_2,Q}^*, \quad \varepsilon \in [0, \infty],$$

if $\mathcal{C}_{L_2,Q}^* < \infty$, and by $\delta_{\max}(\varepsilon, Q) := \infty$ otherwise. Since in the proofs of Lemma 2.9 and Lemma 2.16 we have not used that the inner risks are defined by integrals it is then not hard to see that both lemmas also hold for the above calibration function. Consequently, we say that L_2 is L -calibrated with respect to \mathcal{Q} if

$$\delta_{\max}(\varepsilon, Q) > 0$$

for all $\varepsilon > 0$ and $Q \in \mathcal{Q}$. Analogously, we say that L_2 is uniformly L -calibrated with respect to \mathcal{Q} if $\delta_{\max}(\varepsilon, \mathcal{Q}) := \inf_{Q \in \mathcal{Q}} \delta_{\max}(\varepsilon, Q) > 0$ for all $\varepsilon > 0$. If we now consider a P -instance L_P of L we immediately obtain

$$\delta_{\max, L_P, L_2}(\varepsilon, P(\cdot|x), x) = \delta_{\max, L, L_2}(\varepsilon, P(\cdot|x)), \quad \varepsilon \in [0, \infty], x \in X.$$

In other words, the L -calibration of L_2 can be investigated just like for *regular* losses, i.e. in terms of \mathcal{Q} , while the corresponding results can be used to determine the relation between the excess L_2 -risk and the excess risk of the *unsupervised* loss L_P .

In order to give a first interesting example of template losses let us now turn to our initial question, namely the approximation properties of approximate risk minimizers. To this end let $L : Y \times \mathbb{R} \rightarrow [0, \infty]$ be a regular loss function. We write

$$\mathcal{Q}_{\min}(L) := \{Q : Q \text{ is a distribution on } Y \text{ with } \mathcal{M}_{L,Q}(0^+) \neq \emptyset\}, \tag{25}$$

i.e. $\mathcal{Q}_{\min}(L)$ contains the distributions on Y whose inner L -risks can be minimized. Furthermore, for $Q \in \mathcal{Q}_{\min}(L)$ and $t \in \mathbb{R}$ we define

$$\check{L}(Q, t) := \text{dist}(t, \mathcal{M}_{L,Q}(0^+)) := \inf_{t' \in \mathcal{M}_{L,Q}(0^+)} |t - t'|, \tag{26}$$

i.e. $\check{L}(Q, t)$ measures the distance of t to the set of elements $t' \in \mathbb{R}$ minimizing $\mathcal{C}_{L,Q}(\cdot)$. Now the following fundamental lemma ensures that the above quantity defines a template loss function:

Lemma 2.34 *Let $L : Y \times \mathbb{R} \rightarrow [0, \infty]$ be a regular loss function. Then $\check{L} : \mathcal{Q}_{\min}(L) \times \mathbb{R} \rightarrow [0, \infty)$ defined by (26) is a template loss function.*

Now note that the definition of \check{L} immediately gives $\mathcal{C}_{\check{L},Q}^* = 0$, and therefore we have

$$\mathcal{M}_{\check{L},Q}(\varepsilon) = \{t \in \mathbb{R} : \check{L}(Q, t) < \varepsilon\} = \{t \in \mathbb{R} : \exists t' \in \mathcal{M}_{L,Q}(0^+) \text{ with } |t - t'| < \varepsilon\} \quad (27)$$

for all $Q \in \mathcal{Q}_{\min}(L)$ and $\varepsilon \in [0, \infty]$. Furthermore, we have already mentioned, that the results of Lemma 2.9 remain true for template losses. By Inequality (7) the *self-calibration function* $\delta_{\max, \check{L}, L}(\cdot, Q)$ thus satisfies

$$\delta_{\max, \check{L}, L}(\check{L}(Q, t), Q) \leq \mathcal{C}_{L,Q}(t) - \mathcal{C}_{L,Q}^*$$

for all $Q \in \mathcal{Q}_{\min}(L)$ and all $t \in \mathbb{R}$. Note that if $\mathcal{M}_{L,Q}(0^+)$ contains a single element t_Q^* then the latter inequality becomes

$$\delta_{\max, \check{L}, L}(|t - t_Q^*|, Q) \leq \mathcal{C}_{L,Q}(t) - \mathcal{C}_{L,Q}^*, \quad (28)$$

and consequently, the calibration function quantifies how well an approximate $\mathcal{C}_{L,Q}(\cdot)$ minimizer approximates the exact minimizer t_Q^* . This motivates the following definition:

Definition 2.35 (Self-calibration) *Let $L : Y \times \mathbb{R} \rightarrow [0, \infty]$ be a regular loss function and $\mathcal{Q} \subset \mathcal{Q}_{\min}(L)$. We say that L is (uniformly) self-calibrated with respect to \mathcal{Q} , if L is (uniformly) \check{L} -calibrated with respect to \mathcal{Q} .*

If L is a regular convex loss function then $\mathcal{M}_{L,Q}(0^+)$ is an interval, and hence (27) together with the proof of Lemma 2.11 immediately gives the following result:

Lemma 2.36 *Every regular convex loss function $L : Y \times \mathbb{R} \rightarrow [0, \infty)$ is self-calibrated with respect to $\mathcal{Q}_{\min}(L)$.*

The following proposition shows that for self-calibrated loss functions, approximate risk minimizers approximate the Bayes decision functions:

Theorem 2.37 *Let $L : Y \times \mathbb{R} \rightarrow [0, \infty)$ be a regular loss function that is self-calibrated with respect to some $\mathcal{Q} \subset \mathcal{Q}_{\min}(L)$, and P be a distribution of \mathcal{Q} -type with $\mathcal{R}_{L,P}^* < \infty$. Then for all $\varepsilon > 0$ and $\rho > 0$ there exists a $\delta > 0$ such that for all measurable $f : X \rightarrow \mathbb{R}$ we have*

$$\mathcal{R}_{L,P}(f) < \mathcal{R}_{L,P}^* + \delta \quad \implies \quad P_X(\{x \in X : \check{L}_P(x, f(x)) \geq \rho\}) < \varepsilon.$$

Remark 2.38 In [25] a similar though technically more complicated version of the above theorem has already been proved for binary classification problems.

Remark 2.39 Let L be a convex regular loss such that the sets $\mathcal{M}_{L,P(\cdot|x),x}(0^+)$ are singletons. By Theorem 2.22 there then exists a P_X -almost surely unique minimizer $f_{L,P}^*$ of $\mathcal{R}_{L,P}(\cdot)$, and thus we have $\check{L}_P(x, t) = |t - f_{L,P}^*(x)|$ for P_X -almost all $x \in X$. Moreover, the assumptions of the above theorem are obviously satisfied, and consequently we obtain $f_n \rightarrow f_{L,P}^*$ in probability for all sequences (f_n) with $\mathcal{R}_{L,P}(f_n) \rightarrow \mathcal{R}_{L,P}^*$. In other words, approximate risk minimizers approximate the unique risk minimizer $f_{L,P}^*$ in probability.

Of course, the approximation in probability is a rather weak notion and therefore let us finally describe situations in which we can replace it by a stronger notion of approximation:

Proposition 2.40 *Let $L : Y \times \mathbb{R} \rightarrow [0, \infty)$ be a regular loss function that is self-calibrated with respect to some $\mathcal{Q} \subset \mathcal{Q}_{\min}(L)$, and P be a distribution of \mathcal{Q} -type with $\mathcal{R}_{L,P}^* < \infty$. Assume that there exist $p \in (0, \infty]$, $q > 0$, and a function $b : X \rightarrow [0, \infty]$ with $b^{-1} \in L_p(P_X)$ and*

$$\delta_{\max, \check{L}_P, L}(\varepsilon, P(\cdot|x), x) \geq \varepsilon^q b(x), \quad \varepsilon > 0, x \in X.$$

Then for all measurable $f : X \rightarrow \mathbb{R}$ we have

$$\left(\int_X (\check{L}_P(x, f(x)))^{\frac{pq}{p+1}} dP_X(x) \right)^{\frac{p+1}{pq}} \leq \|b^{-1}\|_{L_p(P_X)}^{\frac{1}{q}} \left(\mathcal{R}_{L,P}(f) - \mathcal{R}_{L,P}^* \right)^{\frac{1}{q}}.$$

Remark 2.41 If $\mathcal{R}_{L,P}(\cdot)$ has an almost surely unique minimizer $f_{L,P}^*$ then Proposition 2.40 ensures $\|\cdot\|_{\frac{pq}{p+1}}$ -convergence of f_n to $f_{L,P}^*$ whenever we have $\mathcal{R}_{L,P}(f_n) \rightarrow \mathcal{R}_{L,P}^*$. Interestingly, if we can only ensure $b^{-1} \in L_{p,\infty}(P_X)$, where $L_{p,\infty}(P_X)$ is a Lorentz space (see e.g. [3]), then the norm $\|\cdot\|_{\frac{pq}{p+1}}$ in the above proposition can be replaced by the Lorentz-norm $\|\cdot\|_{\frac{pq}{p+1}, \infty}$ if we combine Theorem 2.30 with the proof of Theorem 2.37.

3 Examples

In this section we apply the general theory of Section 2 to various loss functions of common learning scenarios such as (weighted) binary classification, regression, density level detection, and density estimation.

3.1 Unweighted Binary Classification

In this section we investigate surrogate loss functions for the binary classification problem. Since such classification calibrated loss functions have already been intensively studied by Bartlett et al. in [1], our main aim in this subsection is to discuss the relationship between their results and our general framework. In addition, the concepts developed in this subsection will be needed in the following subsections when dealing with weighted classification and density level detection.

Let us begin with briefly recalling the unweighted binary classification problem. To this end let $Y := \{-1, 1\}$ and $\mathcal{T} := \mathbb{R}$ throughout the subsection. Our target loss is the regular loss function $L_{\text{class}} : Y \times \mathbb{R} \rightarrow [0, \infty)$ defined by

$$L_{\text{class}}(y, t) := \mathbf{1}_{(-\infty, 0]}(y \operatorname{sign} t), \quad y \in Y, t \in \mathbb{R}, \quad (29)$$

where we use the convention $\operatorname{sign} 0 := 1$. For a given distribution P on $X \times Y$ an easy computation then shows that $\mathcal{R}_{L_{\text{class}}, P}(f)$ is the classification risk discussed in the introduction.

In the following, \mathcal{Q}_Y denotes the set of all distributions on Y . We say that a regular loss function L is (uniformly) *classification calibrated* if it is (uniformly) L_{class} -calibrated with respect to \mathcal{Q}_Y . Furthermore observe that any distribution $Q \in \mathcal{Q}_Y$ can be uniquely described by an $\eta \in [0, 1]$ using the identification $\eta = Q(\{1\})$. If L is a regular loss function we therefore use the notations

$$\mathcal{C}_{L,\eta}(t) := \mathcal{C}_{L,Q}(t), \quad \mathcal{C}_{L,\eta}^* := \mathcal{C}_{L,Q}^*, \quad \mathcal{M}_{L,\eta}(\varepsilon) := \mathcal{M}_{L,Q}(\varepsilon), \quad \text{and} \quad \delta_{\max}(\varepsilon, \eta) := \delta_{\max}(\varepsilon, Q)$$

for $t \in \mathbb{R}$ and $\varepsilon > 0$. Now our first aim is to compute $\mathcal{M}_{L_{\text{class}}, \eta}(\varepsilon)$ and $\delta_{\max}(\varepsilon, \eta)$:

Lemma 3.1 *Let $L : Y \times \mathbb{R} \rightarrow [0, \infty)$ be a regular loss function. Then for all $\eta \in [0, 1]$, $\varepsilon > 0$ we have*

$$\delta_{\max}(\varepsilon, \eta) = \begin{cases} \infty & \text{if } \varepsilon > |2\eta - 1|, \\ \inf_{t \in \mathbb{R}: (2\eta-1) \operatorname{sign} t \leq 0} \mathcal{C}_{L,\eta}(t) - \mathcal{C}_{L,\eta}^* & \text{if } \varepsilon \leq |2\eta - 1|. \end{cases}$$

Proof: For $t \in \mathbb{R}$ a well-known and easy calculation shows

$$\mathcal{C}_{L_{\text{class}},\eta}(t) - \mathcal{C}_{L_{\text{class}},\eta}^* = |2\eta - 1| \cdot \mathbf{1}_{(-\infty,0]}((2\eta - 1) \text{sign } t). \quad (30)$$

Now if $\varepsilon > |2\eta - 1|$ we obviously have $\mathcal{C}_{L_{\text{class}},\eta}(t) - \mathcal{C}_{L_{\text{class}},\eta}^* < \varepsilon$ for all $t \in \mathbb{R}$ and hence we obtain $\mathcal{M}_{L_{\text{class}},\eta}(\varepsilon) = \mathbb{R}$. On the other hand for $\varepsilon \leq |2\eta - 1|$ we have $\mathcal{C}_{L_{\text{class}},\eta}(t) - \mathcal{C}_{L_{\text{class}},\eta}^* < \varepsilon$ if and only if $(2\eta - 1) \text{sign } t > 0$, and hence $\mathcal{M}_{L_{\text{class}},\eta}(\varepsilon) = \{t \in \mathbb{R} : (2\eta - 1) \text{sign } t > 0\}$. ■

In the following we will restrict our considerations to the following class of regular loss functions used in many classification algorithms.

Definition 3.2 *A regular loss function $L : Y \times \mathbb{R} \rightarrow [0, \infty)$ is called margin-based if there exists a $\varphi : \mathbb{R} \rightarrow [0, \infty)$ with $L(y, t) = \varphi(yt)$ for all $y \in Y$ and all $t \in \mathbb{R}$.*

Note that the classification loss L_{class} is *not* margin-based, but many other loss functions are. For some examples we refer to Table 1. The following lemma collects some simple but important properties of margin-based loss functions.

Lemma 3.3 *Let $L : Y \times \mathbb{R} \rightarrow [0, \infty)$ be a margin-based loss function. Then we have:*

- i) For all $\eta \in [0, 1]$ and all $t \in \mathbb{R}$ we have $\mathcal{C}_{L,\eta}(t) = \mathcal{C}_{L,1-\eta}(-t)$ and $\mathcal{C}_{L,\eta}^* = \mathcal{C}_{L,1-\eta}^*$.*
- ii) If L is classification calibrated we have $\mathcal{C}_{L,\eta}(0) > \mathcal{C}_{L,\eta}^*$ for all $\eta \neq 1/2$.*
- iii) L is classification calibrated if and only if the function $H : [0, 1] \rightarrow [0, \infty)$ defined by*

$$H(\eta) := \inf_{t \in \mathbb{R} : (2\eta-1)t \leq 0} \mathcal{C}_{L,\eta}(t) - \mathcal{C}_{L,\eta}^*, \quad \eta \in [0, 1], \quad (31)$$

satisfies $H(\eta) > 0$ for all $\eta \in [0, 1]$ with $\eta \neq 1/2$.

- iv) If L is continuous then we have $\delta_{\max}(\varepsilon, \eta) = H(\eta)$ for all $0 < \varepsilon \leq |2\eta - 1|$.*
- v) The function H is continuous and satisfies $H(\eta) = H(1 - \eta)$, $\eta \in [0, 1]$, and $H(1/2) = 0$.*

Proof: Assertion *i)* is easily verified. Furthermore, for $\eta < 1/2$ assertion *ii)* is guaranteed by $\text{sign } 0 = 1$ and the formula for $\delta_{\max}(\varepsilon, x)$ given in Lemma 3.1. Moreover, for $\eta > 1/2$ the assertion *ii)* follows from $\mathcal{C}_{L,\eta}(0) - \mathcal{C}_{L,\eta}^* = \mathcal{C}_{L,1-\eta}(0) - \mathcal{C}_{L,1-\eta}^* > 0$. Now, *iii)* is a direct consequence of *ii)* and Lemma 3.1. In addition, assertion *iv)* is an easy consequence of $\mathcal{C}_{L,\eta}(0) = \lim_{t \rightarrow 0^-} \mathcal{C}_{L,\eta}(t)$. Finally, the assertions $H(\eta) = H(1 - \eta)$, $\eta \in [0, 1]$, and $H(1/2) = 0$ in *vi)* are trivial and the continuity of H was established in [1, Lem. 5]. ■

Now we can establish the following result which was originally proved by Bartlett, Jordan & McAuliffe in [1].

Theorem 3.4 *Let L be a margin-based loss function. Then the following are equivalent:*

- i) L is classification calibrated.*
- ii) L is uniformly classification calibrated.*

Furthermore, let H be defined by (31). Then for $\delta : [0, 1] \rightarrow [0, \infty)$ defined by $\delta(\varepsilon) := H(\frac{1+\varepsilon}{2})$, $\varepsilon \in [0, 1]$, we have

$$\delta^{**}(\varepsilon) \leq \delta_{\max}^{**}(\varepsilon, \mathcal{Q}_Y), \quad \varepsilon \in [0, 1], \quad (32)$$

and both quantities are actually equal whenever L is continuous.

Proof: We begin with a preliminary consideration. To this end let us fix an $\varepsilon \in (0, 1]$. Then by Lemma 3.1 we observe

$$\inf_{Q \in \mathcal{Q}_Y} \delta_{\max}(\varepsilon, Q) = \inf_{|2\eta-1| \geq \varepsilon} \inf_{t \in \mathbb{R}} \mathcal{C}_{L,\eta}(t) - \mathcal{C}_{L,\eta}^* \geq \inf_{|2\eta-1| \geq \varepsilon} H(\eta) = \inf_{\eta \geq \frac{\varepsilon+1}{2}} H(\eta) =: \tilde{\delta}(\varepsilon). \quad (33)$$

ii) \Rightarrow i). Trivial.

i) \Rightarrow ii). Obviously, H is continuous and strictly positive on all intervals $[\frac{\varepsilon+1}{2}, 1]$, $\varepsilon > 0$, by Lemma 3.3, and consequently we have $\tilde{\delta}(\varepsilon) > 0$ for all $\varepsilon > 0$. Using the monotonicity of $\delta_{\max}(\varepsilon, Q)$ in ε we thus find that L is uniformly classification calibrated.

Finally let us show (32). To this end observe that our preliminary consideration (33) yields $\tilde{\delta}(\varepsilon) \leq \delta_{\max}(\varepsilon, \mathcal{Q}_Y)$ for all $\varepsilon \in [0, 1]$, and consequently we have $\tilde{\delta}^{**}(\varepsilon) \leq \delta_{\max}^{**}(\varepsilon, \mathcal{Q}_Y)$ for all $\varepsilon \in [0, 1]$. Furthermore, we obviously have $\tilde{\delta}(\varepsilon) = \inf_{\varepsilon' \geq \varepsilon} \delta(\varepsilon')$, and hence Lemma A.7 yields $\delta^{**} = \tilde{\delta}^{**}$. In addition, if L is continuous then (33) becomes an equality by Lemma 3.3, and thus repeating the above arguments give $\delta^{**}(\varepsilon) = \delta_{\max}^{**}(\varepsilon, \mathcal{Q}_Y)$ for all $\varepsilon \in [0, 1]$. ■

The above theorem shows that for binary classification, calibration coincides with uniform calibration, and consequently a positive answer to Question 1 automatically gives a positive answer to Question 2. Moreover, in order to find the corresponding inequality, it suffices to compute the function H (or the calibration function for continuous surrogates). For many interesting loss functions this has already been done in [1]. Their results together with the corresponding values of the function δ^{**} defined in Theorem 3.4 are summarized in Table 1. Let us end the discussion with the following two remarks:

Remark 3.5 Let L be a convex, margin-based loss function. Then it was shown in [1] that L is classification calibrated if and only if its associated φ is differentiable at 0 with $\varphi'(0) < 0$. In this case Lemma 2.11 together with the formulas for $\mathcal{M}_{L_{\text{class}}, \eta}(\varepsilon)$ computed in the proof of Lemma 3.1 shows

$$H(\eta) = \varphi(0) - \mathcal{C}_{L,\eta}^*$$

for all $\eta \in [0, 1]$. Moreover, $\eta \mapsto \mathcal{C}_{L,\eta}^*$ is concave, and hence H is convex. Consequently, the formula given in Theorem 3.4 reduces to

$$\delta_{\max}^{**}(\varepsilon, \mathcal{Q}_Y) = \varphi(0) - \mathcal{C}_{L, \frac{1+\varepsilon}{2}}^*, \quad \varepsilon \in [0, 1].$$

Remark 3.6 It is interesting to note that Equation (30) can be used to describe the classification loss by a detection loss function. Indeed, if for a given distribution P on $X \times Y$ we set

$$L_P(x, t) := |2\eta(x) - 1| \cdot \mathbf{1}_{(-\infty, 0]}((2\eta(x) - 1) \text{sign } t), \quad x \in X, t \in \mathbb{R},$$

then $L_P : X \times \mathbb{R} \rightarrow [0, \infty)$ is a detection loss with $h(x) = |2\eta(x) - 1|$, $x \in X$, and (30) states

$$\mathcal{C}_{L_{\text{class}}, \eta(x)}(t) - \mathcal{C}_{L_{\text{class}}, \eta(x)}^* = \mathcal{C}_{L_P, x}(t) - \mathcal{C}_{L_P, x}^*$$

for all $x \in X$, $t \in \mathbb{R}$. Furthermore, Condition (22) is then a weaker version of Tsybakov noise assumption in the sense of [28], and the resulting inequality of Theorem 2.30 essentially coincides with that of [1, Thm. 10]. In this regard it is also interesting to note that the function $\delta^{**}(\varepsilon)$ for the logistic loss (see Table 1) behaves like ε^2 , so that for this loss the inequality of Theorem 2.30 can be substantially simplified.

Sometimes, practical classification problems do not only require a small classification risk but also an estimate of the conditional probability $\eta(x) = P(y = 1|x)$, $x \in X$. If we have a margin-based loss function L for which there is a one-to-one transformation between the sets of minimizers $\mathcal{M}_{L,\eta}(0^+)$ and η then it seems natural to use self-calibration properties of L to investigate whether suitably transformed approximate L -risk minimizers approximate η . This approach is discussed in the following example:

Name of loss function	$\varphi(t)$ for $t \in \mathbb{R}$	$H(\eta)$ for $\eta \in [0, 1]$	$\delta^{**}(\varepsilon)$ for $\varepsilon \in (0, 1]$
Hinge loss	$\max\{0, 1 - t\}$	$ 2\eta - 1 $	ε
Truncated least squares	$(\max\{0, 1 - t\})^2$	$(2\eta - 1)^2$	ε^2
Least squares	$(1 - t)^2$	$(2\eta - 1)^2$	ε^2
Exponential loss	$\exp(-t)$	$1 - 2\sqrt{\eta(1 - \eta)}$	$1 - \sqrt{1 - \varepsilon^2}$
Sigmoid loss	$1 - \tanh(t)$	$ 2\eta - 1 $	ε
Logistic loss	$\ln(1 + e^{-t})$	$\ln 2 + \eta \ln \eta + (1 - \eta) \ln(1 - \eta)$	$\frac{(1 + \varepsilon) \ln(1 + \varepsilon) + (1 - \varepsilon) \ln(1 - \varepsilon)}{2}$

Table 1: Some common margin-based loss functions and the corresponding values for $H(\eta)$ and δ^{**} defined in Lemma 3.3 and Theorem 3.4, respectively. All results are taken from [1] besides the ones for the least squares loss and the logistic loss.

Example 3.7 (Estimating the conditional probabilities with the logistic loss) Let L_{logist} be the logistic loss defined by $\varphi(t) = \ln(1 + e^{-t})$, $t \in \mathbb{R}$. Then it is well-known that

$$\mathcal{M}_{L_{\text{logist}}, \eta}(0^+) = \left\{ \ln\left(\frac{\eta}{1 - \eta}\right) \right\}, \quad \eta \in (0, 1).$$

In other words, if t_η denotes the element contained in $\mathcal{M}_{L_{\text{logist}}, \eta}(0^+)$ then we have $\eta = \frac{1}{1 + e^{-t_\eta}}$. Consequently, if t approximately minimizes $\mathcal{C}_{L_{\text{logist}}, \eta}(\cdot)$ then it is close to t_η by Lemma 2.36 and hence $\frac{1}{1 + e^{-t}}$ can serve as an estimate of η . However, investigating the quality of this estimate by the self-calibration function of L_{logist} causes some technical problems since L_{logist} is only self-calibrated with respect to the distributions $Q \in \mathcal{Q}_Y$ with $Q(\{1\}) \neq 0, 1$. Consequently, we now assess the quality the above estimate *directly*. To this end we define $L : \mathcal{Q}_Y \times \mathbb{R} \rightarrow [0, \infty)$ by $L(\eta, t) := \left| \eta - \frac{1}{1 + e^{-t}} \right|$, $\eta \in [0, 1]$, $t \in \mathbb{R}$. Then L is template loss function which measures the distance between η and its estimate in the sense of the above discussion. Let us compute the calibration function of (L, L_{logist}) . To this end we first observe that $\mathcal{C}_{L, \eta}^* = 0$ for all $\eta \in [0, 1]$, and hence we have $\mathcal{M}_{L, \eta}(\varepsilon) = \{t \in \mathbb{R} : L(\eta, t) < \varepsilon\}$. Consequently, for $\mathcal{C}_\eta(t) := \mathcal{C}_{L_{\text{logist}}, \eta}(t) - \mathcal{C}_{L_{\text{logist}}, \eta}^*$ Lemma 2.11 gives

$$\delta_{\max, L, L_{\text{logist}}}(\varepsilon, \eta) = \min \left\{ \mathcal{C}_\eta \left(\ln \frac{\eta - \varepsilon}{1 - \eta + \varepsilon} \right), \mathcal{C}_\eta \left(\ln \frac{\eta + \varepsilon}{1 - \eta - \varepsilon} \right) \right\},$$

where we use the convention $\mathcal{C}_\eta(\ln -\infty) := \infty$. From this we can easily conclude $\delta_{\max, L, L_{\text{logist}}}(\varepsilon, \eta) = \delta_{\max, L, L_{\text{logist}}}(\varepsilon, 1 - \eta)$ for all $\varepsilon \geq 0$, $\eta \in [0, 1]$. Moreover, some calculations show

$$\eta \ln \frac{\eta}{\eta - \varepsilon} + (1 - \eta) \ln \frac{1 - \eta}{1 - \eta + \varepsilon} \leq \eta \ln \frac{\eta}{\eta + \varepsilon} + (1 - \eta) \ln \frac{1 - \eta}{1 - \eta - \varepsilon}$$

if and only if $\eta \geq \frac{1}{2}$, and consequently for $\eta \in [0, 1/2]$ we find

$$\delta_{\max, L, L_{\text{logist}}}(\varepsilon, \eta) = \begin{cases} \eta \ln \frac{\eta}{\eta + \varepsilon} + (1 - \eta) \ln \frac{1 - \eta}{1 - \eta - \varepsilon} & \text{if } \varepsilon < 1 - \eta \\ \infty & \text{else.} \end{cases}$$

In order to investigate whether L_{logist} is L -calibrated with respect to \mathcal{Q}_Y let us now find a simple lower bound of the above calibration function. To this end let $h(\eta) := \eta \ln \frac{\eta}{\eta + \varepsilon}$. Then its derivative satisfies

$$h'(\eta) = \ln \frac{\eta}{\eta + \varepsilon} + \frac{\varepsilon}{\eta + \varepsilon} = \ln \left(1 - \frac{\varepsilon}{\eta + \varepsilon} \right) + \frac{\varepsilon}{\eta + \varepsilon} \leq -\frac{\varepsilon}{\eta + \varepsilon} + \frac{\varepsilon}{\eta + \varepsilon} = 0,$$

and hence we find $\eta \ln \frac{\eta}{\eta + \varepsilon} \geq \frac{1}{2} \ln \frac{1}{1 + 2\varepsilon}$ for all $\eta \in [0, 1/2]$, $\varepsilon \geq 0$. Analogously, we obtain $(1 - \eta) \ln \frac{1 - \eta}{1 - \eta - \varepsilon} \geq \ln \frac{1}{1 - \varepsilon}$, for $\eta \in [0, 1/2]$, $\varepsilon \in [0, 1 - \eta]$. Both estimates together then yield

$$\delta_{\max, L, L_{\text{logist}}}(\varepsilon, \eta) \geq \frac{1}{2} \ln \frac{1}{1 + 2\varepsilon} + \ln \frac{1}{1 - \varepsilon} \geq \varepsilon^2$$

for all $\eta \in [0, 1/2]$, $\varepsilon \in [0, 1 - \eta]$. From this we easily can conclude that L_{logist} is *uniformly* L -calibrated with respect \mathcal{Q}_Y with $\delta_{\max, L, L_{\text{logist}}}(\varepsilon, \mathcal{Q}_Y) \geq \varepsilon^2$ for all $\varepsilon \geq 0$. If we now consider the squared version L^2 of L then we obviously have $\delta_{\max, L^2, L_{\text{logist}}}(\varepsilon, \mathcal{Q}_Y) = \delta_{\max, L, L_{\text{logist}}}(\sqrt{\varepsilon}, \mathcal{Q}_Y) \geq \varepsilon$ for all $\varepsilon \geq 0$. For measurable $f : X \rightarrow \mathbb{R}$ Theorem 2.13 consequently gives

$$\left(\int_X \left| \eta(x) - \frac{1}{1 + e^{-f(x)}} \right|^2 dP_X(x) \right)^{1/2} \leq \sqrt{\mathcal{R}_{L_{\text{logist}}, P}(f) - \mathcal{R}_{L_{\text{logist}}, P}^*},$$

i.e. the we can assess the quality of the estimate $\frac{1}{1 + e^{-f(x)}}$ in terms of the $\|\cdot\|_2$ -norm.

3.2 Weighted Binary Classification

In this section we investigate surrogate loss functions for weighted binary classification problems. Recall that unlike in the unweighted case, in the weighted binary classification scenario the two types of errors are weighted differently. To make this precise, let $\alpha \in (0, 1)$ be a real number, $Y := \{-1, 1\}$. Then our target loss is the regular loss function $L_{\alpha\text{-class}} : Y \times \mathbb{R} \rightarrow [0, \infty)$ defined by

$$L_{\alpha\text{-class}}(y, t) := \begin{cases} 1 - \alpha & \text{if } y = 1 \text{ and } t < 0 \\ \alpha & \text{if } y = -1 \text{ and } t \geq 0 \\ 0 & \text{otherwise.} \end{cases} \quad (34)$$

Obviously $L_{\alpha\text{-class}}$ is a regular loss function and we have $2L_{\frac{1}{2}\text{-class}} = L_{\text{class}}$. As in the unweighted case we begin with computing the calibration function.

Lemma 3.8 *Let $L : Y \times \mathbb{R} \rightarrow [0, \infty)$ be a regular loss function. Then for all $\alpha \in (0, 1)$, $\eta \in [0, 1]$ and all $\varepsilon > 0$ we have*

$$\delta_{\max}(\varepsilon, \eta) = \begin{cases} \infty & \text{if } \varepsilon > |\eta - \alpha|, \\ \inf_{t \in \mathbb{R} : (\eta - \alpha) \text{ sign } t \leq 0} \mathcal{C}_{L, \eta}(t) - \mathcal{C}_{L, \eta}^* & \text{if } \varepsilon \leq |\eta - \alpha|. \end{cases}$$

Proof: A simple calculation shows $\mathcal{C}_{L_{\alpha\text{-class}}, \eta}(t) - \mathcal{C}_{L_{\alpha\text{-class}}, \eta}^* = |\eta - \alpha| \cdot \mathbf{1}_{(-\infty, 0]}((\eta - \alpha) \text{ sign } t)$, and thus the assertion follows as in the proof of Lemma 3.1. \blacksquare

In the following let L be a margin-based loss function with corresponding function $\varphi : \mathbb{R} \rightarrow [0, \infty)$. For $\alpha \in (0, 1)$ we define the α -weighted version L_α of L by

$$L_\alpha(y, t) := \begin{cases} (1 - \alpha)\varphi(t) & \text{if } y = 1 \\ \alpha\varphi(-t) & \text{if } y = -1, \end{cases} \quad t \in \mathbb{R}.$$

Moreover, when considering these weighted versions, we frequently use the quantities

$$w_\alpha(\eta) := (1 - \alpha)\eta + \alpha(1 - \eta) \quad \text{and} \quad \vartheta_\alpha(\eta) := \frac{(1 - \alpha)\eta}{(1 - \alpha)\eta + \alpha(1 - \eta)}$$

which are defined for all $\eta \in [0, 1]$. Now, the following lemma describes the relation between L_α and L with respect to the inner risks.

Lemma 3.9 *Let L be a margin-based loss function, $\alpha \in (0, 1)$, and $\eta \in [0, 1]$. Then we have:*

- i) $\mathcal{C}_{L_\alpha, \eta}(t) = w_\alpha(\eta) \mathcal{C}_{L, \vartheta_\alpha(\eta)}(t)$ for all $t \in \mathbb{R}$, and $\mathcal{C}_{L_\alpha, \eta}^* = w_\alpha(\eta) \mathcal{C}_{L, \vartheta_\alpha(\eta)}^*$.
- ii) $\min\{\alpha, 1 - \alpha\} \leq w_\alpha(\eta) \leq \max\{\alpha, 1 - \alpha\}$.

iii) If L is classification calibrated and $\eta \neq \alpha$, then $\mathcal{C}_{L,\eta}(0) > \mathcal{C}_{L,\eta}^*$.

Proof: A simple calculation shows $i)$, and $ii)$ easily follows from $w_\alpha(\eta) = (1 - 2\alpha)\eta + \alpha$. Finally, we have $\eta \neq \alpha$ if and only if $\vartheta_\alpha(\eta) \neq \frac{1}{2}$ and hence $iii)$ follows from Lemma 3.3 and

$$\mathcal{C}_{L,\eta}(0) = w_\alpha(\eta) \mathcal{C}_{L,\vartheta_\alpha(\eta)}(0) > w_\alpha(\eta) \mathcal{C}_{L,\vartheta_\alpha(\eta)}^* = \mathcal{C}_{L,\eta}^*.$$

■

With the help of the above lemma we can now characterize when weighted versions of margin based loss functions are $L_{\alpha\text{-class}}$ -calibrated.

Theorem 3.10 *Let L be a margin-based loss function and $\alpha \in (0, 1)$. Then the following statements are equivalent*

- i) L_α is uniformly $L_{\alpha\text{-class}}$ -calibrated with respect to \mathcal{Q}_Y .
- ii) L_α is $L_{\alpha\text{-class}}$ -calibrated with respect to \mathcal{Q}_Y .
- iii) L is classification calibrated.
- iv) The function $H_\alpha : [0, 1] \rightarrow [0, \infty)$ defined by

$$H_\alpha(\eta) := \inf_{t \in \mathbb{R} : (\eta - \alpha)t \leq 0} \mathcal{C}_{L,\eta}(t) - \mathcal{C}_{L,\eta}^*, \quad \eta \in [0, 1], \quad (35)$$

satisfies $H_\alpha(\eta) > 0$ for all $\eta \in [0, 1]$ with $\eta \neq \alpha$.

Furthermore, if one of the statements is true and H is defined by (31) then we have

$$H_\alpha(\eta) = w_\alpha(\eta) H(\vartheta_\alpha(\eta)). \quad (36)$$

Proof: Some easy calculations show $2\vartheta_\alpha(\eta) - 1 = \frac{\eta - \alpha}{(1 - \alpha)\eta + \alpha(1 - \eta)}$ and $|\eta - \alpha| \leq |2\vartheta_\alpha(\eta) - 1|$. Now let $\delta_{\max}(\varepsilon, \eta)$ be defined with respect to L_{class} and L , and $\delta_{\max,\alpha}(\varepsilon, \eta)$ be defined with respect to $L_{\alpha\text{-class}}$ and L_α . Then for $\varepsilon \leq |\eta - \alpha|$ our preliminary considerations yield

$$\begin{aligned} \delta_{\max,\alpha}(\varepsilon, \eta) &= \inf_{\substack{t \in \mathbb{R} \\ (\eta - \alpha) \text{sign } t \leq 0}} \mathcal{C}_{L,\eta}(t) - \mathcal{C}_{L,\eta}^* = w_\alpha(\eta) \inf_{\substack{t \in \mathbb{R} \\ (2\vartheta_\alpha(\eta) - 1) \text{sign } t \leq 0}} \mathcal{C}_{L,\vartheta_\alpha(\eta)}(t) - \mathcal{C}_{L,\vartheta_\alpha(\eta)}^* \\ &= w_\alpha(\eta) \delta_{\max}(\varepsilon, \vartheta_\alpha(\eta)). \end{aligned}$$

From this we immediately obtain the equivalence of $ii)$ and $iii)$. Furthermore, $i) \Rightarrow ii)$ is trivial, and $iii) \Rightarrow i)$ follows from $\delta_{\max,\alpha}(\varepsilon, \eta) = w_\alpha(\eta) \delta_{\max}(\varepsilon, \vartheta_\alpha(\eta)) \geq \min\{\alpha, 1 - \alpha\} \delta_{\max}(\varepsilon, \vartheta_\alpha(\eta))$ and Theorem 3.4. Now, if L is classification calibrated, $iv)$ of Lemma 3.9 together with Lemma 3.8 implies $H_\alpha(\eta) > 0$ for all $\eta \neq \alpha$. Conversely, if $iv)$ holds then we have $\delta_{\max,\alpha}(\varepsilon, \eta) \geq H_\alpha(\eta) > 0$ for $\eta \neq \alpha$ and $0 < \varepsilon \leq |\eta - \alpha|$, and hence L_α is L_{class} -calibrated with respect to \mathcal{Q}_Y . Finally the proof of (36) is analogous to the proof of $\delta_{\max,\alpha}(\varepsilon, \eta) = w_\alpha(\eta) \delta_{\max}(\varepsilon, \vartheta_\alpha(\eta))$. ■

Remark 3.11 The above results can be used to obtain inequalities between the excess risks of $L_{\alpha\text{-class}}$ and α -weighted versions of margin-based, classification calibrated loss functions. In order to find such inequalities let us write

$$\delta_\alpha(\varepsilon) := \inf_{\substack{\eta \in [0, 1] \\ |\eta - \alpha| \geq \varepsilon}} H_\alpha(\eta),$$

Name of loss function	$H(\eta)$	$H_\alpha(\eta)$	$\delta_\alpha^{**}(\varepsilon)$
Hinge loss	$ 2\eta - 1 $	$ \eta - \alpha $	ε
Truncated least squares	$(2\eta - 1)^2$	$\frac{(\eta - \alpha)^2}{\alpha + \eta - 2\alpha\eta}$	$\frac{\varepsilon^2}{2\alpha(1 - \alpha) + \varepsilon(1 - 2\alpha)}$
Least squares	$(2\eta - 1)^2$	$\frac{(\eta - \alpha)^2}{\alpha + \eta - 2\alpha\eta}$	$\frac{\varepsilon^2}{2\alpha(1 - \alpha) + \varepsilon(1 - 2\alpha)}$
Exponential loss	$1 - 2\sqrt{\eta(1 - \eta)}$	$\alpha + \eta - 2\alpha\eta + 2A_0\sqrt{\eta(1 - \eta)}$	$2A_0(A_0 - A_\varepsilon) + \varepsilon(2\alpha - 1)$
Sigmoid loss	$ 2\eta - 1 $	$ \eta - \alpha $	ε

Table 2: The functions H , H_α and δ_α^{**} for some common margin-based losses. For the exponential loss we used the abbreviation $A_\varepsilon := \sqrt{(\alpha - \varepsilon)(1 - \alpha + \varepsilon)}$, $\varepsilon \geq 0$. Furthermore we omitted the logistic loss because the corresponding formulas are too long to fit into the table. Note that for the hinge loss and the sigmoid loss the function δ_α^{**} is actually *independent* of α .

where $\alpha_{\max} := \max\{\alpha, 1 - \alpha\}$ and $\varepsilon \in [0, \alpha_{\max}]$. For $\varepsilon \in [0, \alpha_{\max}]$ Lemma 3.8 then yields

$$\inf_{Q \in \mathcal{Q}_Y} \delta_{\max}(\varepsilon, Q) = \inf_{\substack{\eta \in [0, 1] \\ |\eta - \alpha| \geq \varepsilon}} \inf_{\substack{t \in \mathbb{R} \\ (\eta - \alpha) \operatorname{sign} t \leq 0}} \mathcal{C}_{L_\alpha, \eta}(t) - \mathcal{C}_{L_\alpha, \eta}^* \geq \inf_{\substack{\eta \in [0, 1] \\ |\eta - \alpha| \geq \varepsilon}} H_\alpha(\eta),$$

and consequently we have

$$\delta_\alpha^{**}(\varepsilon) \leq \delta_{\max}^{**}(\varepsilon, \mathcal{Q}_Y).$$

Moreover, it is obvious, that for continuous L , we even have equality in the above formula. Furthermore, for some loss functions we already know $H(\eta)$, $\eta \in [0, 1]$, and hence the computation of $\delta_\alpha^{**}(\varepsilon)$ is straightforward. The corresponding results are summarized in Table 2.

Up to now we have only considered α -weighted versions of classification calibrated loss functions to get $L_{\alpha\text{-class}}$ -calibrated loss functions. We finally show that this is in some sense the only choice:

Proposition 3.12 *Let $\alpha, \beta \in (0, 1)$, L be a margin-based, classification calibrated loss function, and L_β be its β -weighted version. Then L_β is $L_{\alpha\text{-class}}$ -calibrated if and only if $\beta = \alpha$.*

Proof: We have already seen L_α is $L_{\alpha\text{-class}}$ -calibrated, and therefore let us now assume $\alpha \neq \beta$. Without loss of generality we only consider the case $\beta > \alpha$. Then for a fixed $\eta \in (\alpha, \beta)$ an easy computation shows $\vartheta_\beta(\eta) < \frac{1}{2}$, and hence for $\varepsilon > 0$ with $\varepsilon \leq |\eta - \alpha|$ we obtain

$$\delta_{\max}(\varepsilon, \eta) = \inf_{(n - \alpha) \operatorname{sign} t \leq 0} \mathcal{C}_{L_\alpha, \eta}(t) - \mathcal{C}_{L_\alpha, \eta}^* = w_\beta(\eta) \inf_{t < 0} \mathcal{C}_{L, \vartheta_\beta(\eta)}(t) - \mathcal{C}_{L, \vartheta_\beta(\eta)}^*. \quad (37)$$

Furthermore, since L is classification calibrated we have $\inf_{t \geq 0} \mathcal{C}_{L, \vartheta_\beta(\eta)}(t) - \mathcal{C}_{L, \vartheta_\beta(\eta)}^* > 0$, and since $\inf_{t \in \mathbb{R}} \mathcal{C}_{L, \vartheta_\beta(\eta)}(t) - \mathcal{C}_{L, \vartheta_\beta(\eta)}^* = 0$ we find $\inf_{t < 0} \mathcal{C}_{L, \vartheta_\beta(\eta)}(t) - \mathcal{C}_{L, \vartheta_\beta(\eta)}^* = 0$. Together with (37) this shows that L_β is not $L_{\alpha\text{-class}}$ -calibrated. ■

Remark 3.13 Note that the above proposition in particular shows that an α -weighted version of a classification calibrated margin-based loss function is classification calibrated (in the sense of the previous subsection) if and only if $\alpha = \frac{1}{2}$. In other words, changing the weights produces a loss function which is not classification calibrated, and hence this often used heuristic for unbalanced datasets may lead to systematical errors.

3.3 Regression

In this subsection we investigate the use of surrogate loss functions for regression problems. Recall, that the general goal in regression is to predict real-valued outputs, i.e. the label space Y is either the entire real line \mathbb{R} or an interval. The most commonly used loss function for regression is the least squares loss L_{lsquares} defined by $L_{\text{lsquares}}(y, t) := (y - t)^2$. However, this loss is known to be sensitive against outliers, and hence one of our goals is to identify L_{lsquares} -calibrated loss functions. Furthermore, we investigate the problem of finding the mean for symmetric noise distributions and discuss self-calibration for a class of loss functions typically employed in regression problems. Let us begin with introducing this type of loss functions:

Definition 3.14 *Let $L : \mathbb{R} \times \mathbb{R} \rightarrow [0, \infty)$ be a regular loss function. We say that L is:*

- i) distance-based if there is a $\psi : \mathbb{R} \rightarrow [0, \infty)$ with $\psi(0) = 0$ and $L(y, t) = \psi(y - t)$ for all $y, t \in \mathbb{R}$.*
- ii) symmetric if L is distance-based and its associated ψ satisfies $\psi(r) = \psi(-r)$ for all $r \in \mathbb{R}$.*

Obviously, the least squares loss, and more general, the L_p -loss functions, $p > 0$, defined by $|y - t|^p$ are symmetric. Moreover the logistic loss for regression, Huber's loss and the ϵ -insensitive loss are further examples of symmetric loss functions (see also Table 3).

The following definition introduces some important notions for distance-based loss functions:

Definition 3.15 *Let $p > 0$ and $L : \mathbb{R} \times \mathbb{R} \rightarrow [0, \infty)$ be a distance-based loss function with associated $\psi : \mathbb{R} \rightarrow [0, \infty)$. We say that L is:*

- i) (strictly, uniformly) convex if ψ is (strictly, uniformly) convex.*
- ii) of growth type p if there are $c_1, c_2 > 0$ with $c_1 r^p \leq \psi(r) \leq c_2 r^p$ for all sufficiently large r .*
- iii) locally Lipschitz continuous if for all $a > 0$ the restriction $\psi|_{[-a, a]}$ is Lipschitz continuous.*

For our analysis we also need some notions related to distributions on \mathbb{R} . To this end let Q be a distribution on $Y := \mathbb{R}$ with finite first moment, i.e. $|Q|_1 := \mathbb{E}_{y \sim Q} |y| < \infty$. Then the mean of Q is

$$\mathbb{E}Q := \int_{\mathbb{R}} y dQ(y).$$

We call Q *symmetric* around some center $c \in \mathbb{R}$, if $Q(c+A) = Q(c-A)$ for all measurable $A \subset [0, \infty)$. It is not hard to see that the mean $\mathbb{E}Q$ is the only center whenever $|Q|_1 < \infty$. Furthermore we say that Q is symmetric if it is symmetric around some $c \in \mathbb{R}$. Obviously, Q is symmetric around c if and only if its *centered version* $Q^{(c)}$ defined by $Q^{(c)}(A) := Q(c+A)$, $A \subset \mathbb{R}$ measurable, is centered around 0. In the following we denote the set of distributions with first finite moment by $\mathcal{Q}_{\mathbb{R}}$, and the set of distributions with support contained in the *bounded* interval I by \mathcal{Q}_I . Moreover, we write

$$\mathcal{Q}_{\text{bounded}} := \bigcup_{M>0} \mathcal{Q}_{[-M, M]},$$

for the set of distributions with bounded support. Clearly we have $\mathcal{Q}_I \subset \mathcal{Q}_{\text{bounded}} \subset \mathcal{Q}_{\mathbb{R}}$ for all bounded intervals I , and if L is a continuous, distance-based loss function we actually have $\mathcal{Q}_{\text{bounded}} \subset \mathcal{Q}_{\mathbb{R}}(L)$. Moreover, we denote the set of all *symmetric* distributions with first finite moment by \mathcal{Q}_{sym} . Finally, the sets $\mathcal{Q}_{I, \text{sym}}$ for $I \subset \mathbb{R}$ being bounded interval, and $\mathcal{Q}_{\text{bounded}, \text{sym}}$ are defined in the obvious way.

Our first goal is to identify distance-based, L_{lsquares} -calibrated loss functions L . To this end recall

$$\mathcal{M}_{L_{\text{lsquares}}, Q}(0^+) = \{\mathbb{E}Q\}$$

for all $Q \in \mathcal{Q}(L_{\text{lsquares}})$, and consequently, if L is a L_{lsquares} -calibrated loss function then we must have $\mathcal{M}_{L, Q}(0^+) \subset \{\mathbb{E}Q\}$ for all $Q \in \mathcal{Q}(L)$. This observation motivates the following two propositions in which we investigate the sets $\mathcal{M}_{L, Q}(0^+)$ for distance-based loss functions.

Proposition 3.16 *Let L be a convex, distance-based loss function whose associated ψ satisfies $\lim_{r \rightarrow \pm\infty} \psi(r) = \infty$, and let $Q \in \mathcal{Q}_{\mathbb{R}}$ with $\mathcal{C}_{L, Q}(t) < \infty$ for all $t \in \mathbb{R}$. Then $\mathcal{M}_{L, Q}(0^+)$ is a compact, non-empty interval. Moreover, if ψ is strictly convex then $\mathcal{M}_{L, Q}(0^+)$ contains exactly one element.*

Proof: Our first goal is to show that $\lim_{t \rightarrow \pm\infty} \mathcal{C}_{L, Q}(t) = \infty$. To this end let $(t_n) \subset \mathbb{R}$ be a sequence with $t_n \rightarrow -\infty$, and $B > 0$. Since $\lim_{r \rightarrow \pm\infty} \psi(r) = \infty$ there then exists an $r_0 > 0$ such that $\psi(r) \geq 2B$ for all $r \in \mathbb{R}$ with $|r| \geq r_0$. Since $Q(\mathbb{R}) = 1$ there is also an $M > 0$ with $Q([-M, M]) \geq 1/2$. Finally, there exists an $n_0 \geq 1$ with $t_n \leq -M - r_0$ for all $n \geq n_0$. For $y \in [-M, M]$ this yields $y - t_n \geq r_0$, and hence we find $\psi(y - t_n) \geq 2B$ for all $n \geq n_0$. From this we easily find

$$\mathcal{C}_{L, Q}(t_n) \geq \int_{[-M, M]} \psi(y - t_n) dQ(y) \geq 2B Q([-M, M]) = B,$$

i.e. we have shown $\mathcal{C}_{L, Q}(t_n) \rightarrow \infty$. Analogously we can show $\lim_{t \rightarrow \infty} \mathcal{C}_{L, Q}(t) = \infty$, and consequently we have $\lim_{t \rightarrow \pm\infty} \mathcal{C}_{L, Q}(t) = \infty$. Furthermore, the convexity of ψ implies that $t \mapsto \mathcal{C}_{L, Q}(t)$ is convex and hence it is continuous by the assumption $\mathcal{C}_{L, Q}(t) < \infty$, $t \in \mathbb{R}$. Now the assertions follow. ■

Note that for distributions $Q \in \mathcal{Q}_{\text{bounded}}$ we automatically have $\mathcal{C}_{L, Q}(t) < \infty$ for all $t \in \mathbb{R}$. Furthermore, if L is of some growth type p then we have $\mathcal{C}_{L, Q}(t) < \infty$ for all $t \in \mathbb{R}$ and all $Q \in \mathcal{Q}_{\mathbb{R}}(L)$. Consequently, the above proposition gives $\mathcal{M}_{L, Q}(0^+) \neq \emptyset$ in both cases.

The following proposition compares $\mathcal{M}_{L, Q}(0^+)$ with the mean $\mathbb{E}Q$. Note that a similar result was independently found by A. Caponnetto in [6].

Proposition 3.17 *Let $M > 0$, and L be a distance-based, locally Lipschitz continuous loss function with associated ψ . Then we have:*

- i) *If $\mathbb{E}Q \in \mathcal{M}_{L, Q}(0^+)$ for all $Q \in \mathcal{Q}_{[-M, M], \text{sym}}$, then L is symmetric.*
- ii) *If $\mathbb{E}Q \in \mathcal{M}_{L, Q}(0^+)$ for all $Q \in \mathcal{Q}_{\text{bounded}}$, then there exists a $c > 0$ with $L = c L_{\text{lsquares}}$.*

Proof: Recall that the fundamental theorem of calculus for Lebesgue integrals shows that the derivative ψ' is (Lebesgue)-almost surely defined and integrable on every bounded interval.

i). Let us fix a $y \in [-M, M]$ such that ψ is differentiable at y and $-y$. We define $Q := \frac{1}{2}\delta_{\{-y\}} + \frac{1}{2}\delta_{\{y\}}$. Then we have $Q \in \mathcal{Q}_{\text{sym}}$ with $\mathbb{E}Q = 0$, and $\mathcal{C}_{L, Q}(t) = \frac{1}{2}\psi(-y-t) + \frac{1}{2}\psi(y-t)$. Therefore the derivative of $\mathcal{C}_{L, Q}(\cdot)$ exists at 0 and can be computed by $\mathcal{C}'_{L, Q}(0) = -\frac{1}{2}\psi'(-y) + \frac{1}{2}\psi'(y)$. Furthermore, our assumption shows that $\mathcal{C}_{L, Q}(\cdot)$ has a minimum at 0, and hence we have $0 = \mathcal{C}'_{L, Q}(0)$, i.e. $\psi'(-y) = -\psi'(y)$. According to our preliminary remark, this equality holds for almost all y , and thus the fundamental theorem of calculus for Lebesgue integrals shows that for all $y_0 \in \mathbb{R}$ we have

$$\psi(y_0) = \psi(0) + \int_0^{y_0} \psi'(t) dt = \psi(0) - \int_0^{y_0} \psi'(-t) dt = \psi(0) - \int_{-y_0}^0 \psi'(t) dt = \psi(-y_0).$$

ii). Let $y \neq 0$ and $\alpha > 0$ be real numbers such that ψ is differentiable at y , $-y$, and αy . We define $Q := \frac{\alpha}{1+\alpha}\delta_{\{0\}} + \frac{1}{1+\alpha}\delta_{\{(1+\alpha)y\}}$, so that we obtain $\mathbb{E}Q = y$ and $\mathcal{C}_{L,Q}(t) = \frac{\alpha}{1+\alpha}\psi(-t) + \frac{1}{1+\alpha}\psi(y + \alpha y - t)$ for all $t \in \mathbb{R}$. This shows that the derivative of $\mathcal{C}_{L,Q}(\cdot)$ exists at y and can be computed by

$$\mathcal{C}'_{L,Q}(y) = -\frac{\alpha}{1+\alpha}\psi'(-y) - \frac{1}{1+\alpha}\psi'(\alpha y) = \frac{\alpha}{1+\alpha}\psi'(y) - \frac{1}{1+\alpha}\psi'(\alpha y),$$

where in the last step we used i). Again, our assumption $y \in \mathcal{M}_{L,Q}(0^+)$ gives $\mathcal{C}'_{L,Q}(y) = 0$ and hence we find $\alpha\psi'(y) = \psi'(\alpha y)$. Obviously, the latter holds for almost all $\alpha > 0$ and thus we obtain

$$\psi(ty) = \psi(0) + \int_0^t \psi'(sy)y ds = \int_0^t s\psi'(y)y ds = \frac{\psi'(y)}{2y}(ty)^2$$

for all $t > 0$. From this we easily obtain the assertion. ■

The above proposition shows that there are only trivial surrogates for the least squares loss if one is interested in the entire class $\mathcal{Q}(L_{\text{lsquares}}) = \{Q \in \mathcal{Q}_{\mathbb{R}} : \mathbb{E}_{y \sim Q} y^2 < \infty\}$. Furthermore, it shows that the least squares loss is essentially the only loss function whose minimizer is the mean for all distributions in $\mathcal{Q}(L_{\text{lsquares}})$, i.e. if we are actually interested in finding the regression function $x \mapsto \mathbb{E}_P(Y|x)$, and we just know $\mathbb{E}_{x \sim P_X} \mathbb{E}_{y \sim P(\cdot|x)} y^2 < \infty$, then the least squares loss is the *only* loss for this task. However, if we cannot ensure this tail assumption but know that the conditional distributions $P(\cdot|x)$ are *symmetric* then the above proposition suggests that we may actually have alternatives to the least squares loss. In order to investigate this conjecture systematically we first need a target loss function that describes the goal of finding the mean. To this end let us consider the template loss $L_{\text{mean}} : \mathcal{Q}_{\mathbb{R}} \times \mathbb{R} \rightarrow [0, \infty)$ defined by

$$L_{\text{mean}}(Q, t) := |\mathbb{E}Q - t|, \quad t \in \mathbb{R}, Q \in \mathcal{Q}_{\mathbb{R}}.$$

Note that this indeed defines a template loss since given a distribution P on $X \times \mathbb{R}$ of $\mathcal{Q}_{\mathbb{R}}$ -type it is easy to see that $(x, t) \mapsto L_{\text{mean}}(P(\cdot|x), t) = |\mathbb{E}_P(Y|x) - t|$ is measurable. Moreover we have

$$L_{\text{mean}}^2(Q, t) = (\mathbb{E}Q - t)^2 = \mathcal{C}_{L_{\text{lsquares}}, Q}(t) - \mathcal{C}_{L_{\text{lsquares}}, Q}^*$$

for all $Q \in \mathcal{Q}(L_{\text{lsquares}})$ and $t \in \mathbb{R}$. Since the minimal L_{mean} -risks equal 0, i.e. $\mathcal{C}_{L_{\text{mean}}, Q}^* = 0$, we consequently have $\mathcal{M}_{L_{\text{mean}}, Q}(\sqrt{\varepsilon}) = \mathcal{M}_{L_{\text{lsquares}}, Q}(\varepsilon)$ for all $\varepsilon > 0$, and from this we obtain

$$\delta_{\max, L_{\text{mean}}, L}(\sqrt{\varepsilon}, Q) = \delta_{\max, L_{\text{lsquares}}, L}(\varepsilon, Q) \quad (38)$$

for all distance-based losses L , all $Q \in \mathcal{Q}(L_{\text{lsquares}}) \cap \mathcal{Q}(L)$, and all $\varepsilon \in [0, \infty]$. In other words, by considering L_{mean} -calibration we simultaneously obtain results on L_{lsquares} -calibration. Furthermore, if $t \mapsto \mathcal{C}_{L,Q}(t)$ has a *unique* minimum at $\mathbb{E}Q$ then we obviously have $L_{\text{mean}}(Q, t) = \check{L}(Q, t)$ for all $t \in \mathbb{R}$, and consequently we obtain

$$\delta_{\max, L_{\text{mean}}, L}(\varepsilon, Q) = \delta_{\max, \check{L}, L}(\varepsilon, Q), \quad \varepsilon \in [0, \infty]. \quad (39)$$

In other words, by considering L_{mean} -calibration of L we will also gain some insight into the self-calibration properties of L .

Now, the following key lemma presents an alternative way to compute the inner risks $\mathcal{C}_{L,Q}(\cdot)$ when both L and Q are symmetric:

Lemma 3.18 *Let L be a symmetric loss function with associated ψ , and $Q \in \mathcal{Q}_{\text{sym}}(L)$. For $t \in \mathbb{R}$ we then have*

$$\mathcal{C}_{L,Q}(\mathbb{E}Q + t) = \mathcal{C}_{L,Q}(\mathbb{E}Q - t) = \frac{1}{2} \int_{\mathbb{R}} (\psi(y - \mathbb{E}Q - t) + \psi(y - \mathbb{E}Q + t)) dQ(y). \quad (40)$$

In addition, if L is convex we have $\mathcal{C}_{L,Q}(\mathbb{E}Q) = \mathcal{C}_{L,Q}^$, and if L is actually strictly convex we also have $\mathcal{C}_{L,Q}(\mathbb{E}Q + t) > \mathcal{C}_{L,Q}^*$ for all $t \neq 0$.*

Proof: Let us fix a $Q \in \mathcal{Q}_{\text{sym}}(L)$ and write $m := \mathbb{E}Q$. The symmetry of $Q^{(m)}$ and ψ then yields

$$\mathcal{C}_{L,Q}(m + t) = \int_{\mathbb{R}} \psi(y - t) dQ^{(m)}(y) = \int_{\mathbb{R}} \psi(-y - t) dQ^{(m)}(y) = \int_{\mathbb{R}} \psi(y + t) dQ^{(m)}(y) = \mathcal{C}_{L,Q}(m - t).$$

Since this gives $\mathcal{C}_{L,Q}(m + t) = \frac{1}{2}(\mathcal{C}_{L,Q}(m + t) + \mathcal{C}_{L,Q}(m - t))$ we also obtain the second equation. Furthermore, if ψ is convex we can easily conclude

$$\mathcal{C}_{L,Q}(m + t) = \frac{1}{2} \int_{\mathbb{R}} (\psi(y - t) + \psi(y + t)) dQ^{(m)}(y) \geq \int_{\mathbb{R}} \psi(y) dQ^{(m)}(y) = \mathcal{C}_{L,Q}(m)$$

for all $t \in \mathbb{R}$. This shows the second assertion. The third assertion can be shown analogously. ■

With the help of the above lemma we can derive a simple formula for the calibration function $\delta_{\max, L_{\text{mean}}, L}(\varepsilon, Q)$ if L is convex:

Lemma 3.19 *Let L be a symmetric, convex loss function and $Q \in \mathcal{Q}_{\text{sym}}(L)$. Then for all $\varepsilon \geq 0$ we have*

$$\delta_{\max, L_{\text{mean}}, L}(\varepsilon, Q) = \mathcal{C}_{L,Q}(\mathbb{E}Q + \varepsilon) - \mathcal{C}_{L,Q}^*. \quad (41)$$

Proof: Obviously, we have $\mathcal{M}_{L_{\text{mean}}, Q}(\varepsilon) = (\mathbb{E}Q - \varepsilon, \mathbb{E}Q + \varepsilon)$. Since this set is open, it is easy to see that the continuity assumption in Lemma 3.18 is superfluous and hence the assertion follows by Lemma 2.11 and Lemma 3.18. ■

Our next result is a rather technical lemma will be used at various times to establish *upper* bounds on $\delta_{\max}(\varepsilon, Q)$. For its formulation we need the set

$$\mathcal{Q}_{\text{sym}}^* := \{Q \in \mathcal{Q}_{\text{sym}} : Q(\mathbb{E}Q + (-\rho, \rho)) > 0 \text{ for all } \rho > 0\}$$

which contains all symmetric distributions that do not vanish around their means. Moreover, we also need the sets $\mathcal{Q}_{I, \text{sym}}^* := \mathcal{Q}_I \cap \mathcal{Q}_{\text{sym}}^*$ for I being a bounded interval, and $\mathcal{Q}_{\text{bounded}, \text{sym}}^* := \mathcal{Q}_{\text{bounded}} \cap \mathcal{Q}_{\text{sym}}^*$. Now the result reads as follows:

Lemma 3.20 *Let L be a symmetric, continuous loss function with associated ψ . Assume that there exist $\delta_0 \in \mathbb{R}$, $s_1, s_2 \in \mathbb{R}$, $s_1 \neq s_2$ and an $\varepsilon_0 > 0$ such that for all $\varepsilon \in [0, \varepsilon_0]$ we have*

$$\frac{\psi(s_1 + \varepsilon) + \psi(s_2 + \varepsilon)}{2} - \psi\left(\frac{s_1 + s_2}{2} + \varepsilon\right) \leq \delta_0. \quad (42)$$

Then for $M := \left\lceil \frac{s_1 + s_2}{2} \right\rceil + \varepsilon_0$ and all $\delta > 0$ there exists a Lebesgue-absolutely continuous $Q \in \mathcal{Q}_{[-M, M], \text{sym}}^(L)$ with $\mathbb{E}Q = 0$ such that for $t := \frac{s_2 - s_1}{2}$ we have*

$$\mathcal{C}_{L,Q}(t) - \mathcal{C}_{L,Q}(0) \leq \delta_0 + \delta.$$

Furthermore, there exists a Lebesgue-absolutely continuous $Q \in \mathcal{Q}_{[-M, M], \text{sym}}(L)$ with $\mathbb{E}Q = 0$ and

$$\mathcal{C}_{L,Q}(t) - \mathcal{C}_{L,Q}(0) \leq \delta_0.$$

Proof: We write $y_0 := \lfloor \frac{s_1+s_2}{2} \rfloor$. Furthermore, if $y_0 = 0$ we let Q be the uniform distribution $\mu_{[-\varepsilon_0, \varepsilon_0]}$ on $[-\varepsilon_0, \varepsilon_0]$. Otherwise, we define

$$Q := \alpha \mu_{[-\frac{y_0}{2}, \frac{y_0}{2}]} + \frac{1-\alpha}{2} \mu_{[-y_0-\varepsilon_0, -y_0]} + \frac{1-\alpha}{2} \mu_{[y_0, y_0+\varepsilon_0]},$$

where $\alpha \in (0, 1)$ is a real number satisfying

$$\sup_{y \in [-\frac{y_0}{2}, \frac{y_0}{2}]} \left| \frac{\psi(y-t) + \psi(y+t)}{2} - \psi(y) \right| \leq \frac{\delta}{\alpha}.$$

Now, in the case $y_0 \neq 0$ this construction yields $\mathbb{E}Q = 0$ and

$$\begin{aligned} \mathcal{C}_{L,Q}(t) - \mathcal{C}_{L,Q}(0) &= \int_{\mathbb{R}} \frac{\psi(y-t) + \psi(y+t)}{2} - \psi(y) dQ(y) \\ &= \alpha \int_{[-\frac{y_0}{2}, \frac{y_0}{2}]} \frac{\psi(y-t) + \psi(y+t)}{2} - \psi(y) \mu_{[-\frac{y_0}{2}, \frac{y_0}{2}]}(y) \\ &\quad + (1-\alpha) \int_{[y_0, y_0+\varepsilon_0]} \frac{\psi(y-t) + \psi(y+t)}{2} - \psi(y) \mu_{[y_0, y_0+\varepsilon_0]}(y) \\ &\leq \delta + (1-\alpha) \int_{[0, \varepsilon_0]} \frac{\psi(s_1+\varepsilon) + \psi(s_2+\varepsilon)}{2} - \psi\left(\frac{s_1+s_2}{2} + \varepsilon\right) \mu_{[0, \varepsilon_0]}(y) \\ &\leq \delta_0 + \delta. \end{aligned}$$

Furthermore, the case $y_0 = 0$ can be shown analogously, and the last assertion follows if we repeat the above construction with $\alpha = 0$. ■

With the above preparations we can finally establish our first main result that characterizes loss functions L that are L_{mean} -calibrated with respect to $\mathcal{Q}_{\text{sym}}^*(L)$:

Theorem 3.21 *Let L be a symmetric, continuous loss function. Then the following statements are equivalent:*

- i) L is L_{mean} -calibrated with respect to $\mathcal{Q}_{\text{sym}}^*(L)$.
- ii) L is convex and its associated function ψ has its only minimum at 0.

Proof: $ii) \Rightarrow i)$. Assume that L is not L_{mean} -calibrated with respect to $\mathcal{Q}_{\text{sym}}^*(L)$. By Lemma 3.19 there then exist a $Q \in \mathcal{Q}_{\text{sym}}^*(L)$ and a $t \neq 0$ with $\mathcal{C}_{L,Q}(m+t) = \mathcal{C}_{L,Q}^*$, where $m := \mathbb{E}Q$. Using $\mathcal{C}_{L,Q}(m) = \mathcal{C}_{L,Q}^*$, which we know from Lemma 3.18, then yields

$$\int_{\mathbb{R}} \frac{\psi(y-t) + \psi(y+t)}{2} - \psi(y) dQ^{(m)}(y) = \mathcal{C}_{L,Q}(m+t) - \mathcal{C}_{L,Q}(m) = 0,$$

and hence the convexity of ψ shows $\frac{\psi(y-m-t) + \psi(y-m+t)}{2} - \psi(y-m) = 0$ for Q -almost all $y \in \mathbb{R}$. The continuity of ψ and the assumption $Q(m + (-\rho, \rho)) > 0$ for all $\rho > 0$, then guarantee that $\frac{\psi(y-m-t) + \psi(y-m+t)}{2} - \psi(y-m) = 0$ holds for $y := m$. However, by the symmetry of ψ this implies $\psi(t) = \psi(0)$ which violates our assumption on ψ .

$i) \Rightarrow ii)$. Assume that ψ is not convex. Then Lemma A.1 shows that there exist $s_1, s_2 \in \mathbb{R}$ with $\frac{\psi(s_1) + \psi(s_2)}{2} - \psi(\frac{s_1+s_2}{2}) < 0$. With the continuity of ψ we then find (42) for some suitable $\delta_0 < 0$ and $\varepsilon_0 > 0$, and consequently Lemma 3.20 gives a $Q \in \mathcal{Q}_{[-M, M], \text{sym}}^*(L)$ and a $t \neq 0$ with $\mathbb{E}Q = 0$ and

$\mathcal{C}_{L,Q}(t) < \mathcal{C}_{L,Q}(0)$. Now observe that since ψ is continuous and Q has bounded support, the map $t \mapsto \mathcal{C}_{L,Q}(t)$ is continuous on \mathbb{R} . Let $(t_n) \subset \mathbb{R}$ be a sequence with $\mathcal{C}_{L,Q}(t_n) \rightarrow \mathcal{C}_{L,Q}^*$ for $n \rightarrow \infty$. Now, our previous considerations showed $\mathcal{C}_{L,Q}(0) \neq \mathcal{C}_{L,Q}^*$ and hence (t_n) is eventually bounded away from 0, i.e. there exist an $\varepsilon > 0$ and an $n_0 \in \mathbb{N}$ such that $|t_n| \geq \varepsilon$ for all $n \geq n_0$. This gives

$$\delta_{\max}(\varepsilon, Q) = \inf_{t' \notin (-\varepsilon, \varepsilon)} \mathcal{C}_{L,Q}(t') - \mathcal{C}_{L,Q}^* \leq \mathcal{C}_{L,Q}(t_n) - \mathcal{C}_{L,Q}^*,$$

for all $n \geq n_0$. For $n \rightarrow \infty$ we hence find $\delta_{\max}(\varepsilon, Q) = 0$, and consequently L is convex. Finally, assume that there exists a $t \neq 0$ with $\psi(t) = \psi(0)$. Then we find $\mathcal{C}_{L,Q}(t) = \mathcal{C}_{L,Q}^*$ for the distribution Q defined by $Q(\{0\}) = 1$, and hence we obtain $\delta_{\max}(t, Q) = 0$ by Lemma 3.19. ■

The following theorem considers calibration with respect to the bigger class \mathcal{Q}_{sym} :

Theorem 3.22 *Let L be a symmetric, continuous loss function. Then the following statements are equivalent:*

- i) L is L_{mean} -calibrated with respect to $\mathcal{Q}_{\text{sym}}(L)$.
- ii) L is strictly convex.

Proof: If L is strictly convex then Lemma 3.18 and Lemma 3.19 show that L is L_{mean} -calibrated with respect to $\mathcal{Q}_{\text{sym}}(L)$. Conversely, if L is L_{mean} -calibrated with respect to $\mathcal{Q}_{\text{sym}}(L)$, then Theorem 3.21 shows that L is convex. Let us suppose that its associated $\psi : \mathbb{R} \rightarrow [0, \infty)$ is not strictly convex. Then there exists $s_1, s_2 \in \mathbb{R}$ with $s_1 \neq s_2$ and

$$\psi\left(\frac{1}{2}s_1 + \frac{1}{2}s_2\right) = \frac{1}{2}\psi(s_1) + \frac{1}{2}\psi(s_2).$$

From this and Lemma A.1 we easily find (42) for $\delta_0 = 0$ and some suitable $\varepsilon_0 > 0$. Lemma 3.20 then gives a $t_0 \neq 0$ and a $Q \in \mathcal{Q}_{[-M, M], \text{sym}}(L)$ with $\mathcal{C}_{L,Q}(\mathbb{E}Q + t_0) = \mathcal{C}_{L,Q}^*$ and hence Lemma 3.19 shows that L is not L_{mean} -calibrated with respect to $\mathcal{Q}_{\text{sym}}(L)$. ■

Our next aim is to estimate the function $\varepsilon \mapsto \delta_{\max}(\varepsilon, Q)$ for some classes $Q \subset \mathcal{Q}_{\text{sym}}$. To this end we define the *modulus of convexity* of a function $f : I \rightarrow \mathbb{R}$, where I is an interval, by

$$\delta_f(\varepsilon) := \inf \left\{ \frac{f(x_1) + f(x_2)}{2} - f\left(\frac{x_1 + x_2}{2}\right) : x_1, x_2 \in I \text{ with } |x_1 - x_2| \geq \varepsilon \right\}, \quad \varepsilon > 0.$$

In addition we say that f is *uniformly convex* if $\delta_f(\varepsilon) > 0$ for all $\varepsilon > 0$. Some properties of the modulus of convexity and uniformly convex functions can be found in Appendix.

With the help of the modulus of convexity we can now formulate the following theorem that estimates $\delta_{\max}(\varepsilon, Q)$ and characterizes uniform calibration:

Theorem 3.23 *Let L be a symmetric, convex loss function with associated ψ . Then we have:*

- i) For all $M > 0$, $\varepsilon > 0$, and all Q with $\mathcal{Q}_{[-M, M], \text{sym}}^* \subset Q \subset \mathcal{Q}_{[-M, M], \text{sym}}$ we have

$$\delta_{\psi|_{[-(2M+\varepsilon), 2M+\varepsilon]}}(2\varepsilon) \leq \delta_{\max}(\varepsilon, Q) \leq \delta_{\psi|_{[-M/2, M/2]}}(2\varepsilon). \quad (43)$$

Consequently, the following statements are equivalent:

- (a) L is uniformly L_{mean} -calibrated with respect to $\mathcal{Q}_{[-M, M], \text{sym}}^*$ for all $M > 0$.

(b) L is uniformly L_{mean} -calibrated with respect to $\mathcal{Q}_{[-M,M],\text{sym}}$ for all $M > 0$.

(c) The function ψ is strictly convex.

ii) For all $\varepsilon > 0$ we have

$$\delta_\psi(2\varepsilon) = \delta_{\max}(\varepsilon, \mathcal{Q}_{\text{sym}}(L)) = \delta_{\max}(\varepsilon, \mathcal{Q}_{\text{bounded,sym}}^*(L)). \quad (44)$$

Consequently, the following statements are equivalent:

(a) L is uniformly L_{mean} -calibrated with respect to $\mathcal{Q}_{\text{sym}}(L)$.

(b) L is uniformly L_{mean} -calibrated with respect to $\mathcal{Q}_{\text{bounded,sym}}^*(L)$.

(c) The function ψ is uniformly convex.

Proof: i). For $Q \in \mathcal{Q}_{[-M,M],\text{sym}}$ we have $\mathbb{E}Q \in [-M, M]$, and hence we find

$$\delta_{\max}(\varepsilon, Q) = \int_{[-M,M]} \frac{\psi(y - \mathbb{E}Q - \varepsilon) + \psi(y - \mathbb{E}Q + \varepsilon)}{2} - \psi(y - \mathbb{E}Q) dQ(y) \geq \delta_{\psi|[-(2M+\varepsilon), 2M+\varepsilon]}(2\varepsilon)$$

by Lemma 3.18 and Lemma 3.19. This shows the first inequality. To prove the second inequality let $n \geq 1$. Then there exist $s_1, s_2 \in [-M/2, M/2]$ with $s_1 - s_2 \geq 2\varepsilon$ and

$$\frac{\psi(s_1) + \psi(s_2)}{2} - \psi\left(\frac{s_1 + s_2}{2}\right) < \delta_{\psi|[-M/2, M/2]}(2\varepsilon) + \frac{1}{n} =: \delta_0.$$

By the continuity of ψ there thus exists an $\varepsilon_0 \in (0, M]$ such that (42) is satisfied for δ_0 , and consequently Lemma 3.20 gives a $Q \in \mathcal{Q}_{[-M,M],\text{sym}}^*$ with $\mathbb{E}Q = 0$ and with

$$\mathcal{C}_{L,Q}(t) - \mathcal{C}_{L,Q}^* \leq \delta_{\psi|[-M/2, M/2]}(2\varepsilon) + \frac{2}{n}$$

for $t := \frac{s_1 - s_2}{2}$. Since $t \geq \varepsilon$ we also have $\delta_{\max}(\varepsilon, Q) \leq \mathcal{C}_{L,Q}(t) - \mathcal{C}_{L,Q}^*$, and hence we find

$$\delta_{\max}(\varepsilon, \mathcal{Q}_{[-M,M],\text{sym}}^*) \leq \delta_{\max}(\varepsilon, Q) \leq \delta_{\psi|[-M/2, M/2]}(2\varepsilon) + \frac{2}{n}.$$

Since this holds for all $n \geq 1$ the second inequality follows. Finally, Lemma A.1 shows that ψ is strictly convex if and only if $\delta_{\psi|[-B,B]}(\varepsilon) > 0$ for all $B, \varepsilon > 0$, and hence the characterization follows.

ii). Analogously to the proof of the first inequality in (43) we find $\delta_\psi(2\varepsilon) \leq \delta_{\max}(\varepsilon, \mathcal{Q}_{\text{sym}}(L))$ for all $\varepsilon > 0$. Furthermore by the second inequality in (43) we obtain $\delta_{\max}(\varepsilon, \mathcal{Q}_{\text{bounded,sym}}^*(L)) \leq \delta_\psi(2\varepsilon)$ for all $\varepsilon > 0$, and hence (44) is proved. Finally, the characterization is a consequence of (44). ■

Remark 3.24 The above theorem shows that the modulus of convexity completely determines whether a loss function is uniformly L_{mean} -calibrated with respect to $\mathcal{Q}_{\text{sym}}(L)$ or $\mathcal{Q}_{\text{bounded,sym}}^*(L)$. Unfortunately, Lemma A.3 shows that for all distance-based loss functions of growth type $p < 2$ we have $\delta_\psi(\varepsilon) = 0$ for all $\varepsilon > 0$. In particular, Lipschitz continuous, distance-based losses which are of special interest for robust regression methods (see e.g. [8]) are not uniformly calibrated with respect to $\mathcal{Q}_{\text{sym}}(L)$ or $\mathcal{Q}_{\text{bounded,sym}}^*(L)$, and consequently we cannot establish *strong* relations between the excess L -risks and $\mathcal{R}_{L_{\text{mean}},P}(\cdot)$ in the sense of Question 2.

On the other hand, note that symmetric, strictly convex losses L are L_{mean} -calibrated with respect to $\mathcal{Q}_{\text{sym}}(L)$, and hence we can show analogously to Theorem 2.37 that $f_n \rightarrow \mathbb{E}(Y|\cdot)$ in probability, whenever $\mathcal{R}_{L,P}(f_n) \rightarrow \mathcal{R}_{L,P}^*$ and P is of type $\mathcal{Q}_{\text{sym}}(L)$. In addition, if we restrict our considerations to $\mathcal{Q}_{[-M,M],\text{sym}}$ or $\mathcal{Q}_{[-M,M],\text{sym}}^*$ then every strictly convex loss becomes uniformly L_{mean} -calibrated, and in this case $\delta_{\psi|[-B,B]}(\cdot)$, $B > 0$, can be used to describe the corresponding calibration function. For some important losses we have listed the behaviour of $\delta_{\psi|[-B,B]}(\cdot)$ in Table 3. Furthermore, Lemma A.3 establishes a formula for the modulus of convexity which often helps to bound the modulus.

Loss function	$\psi(t)$	lower bound of $\delta_{\psi [-B,B]}$	upper bound of $\delta_{\psi [-B,B]}$
L_1 -loss	$ t $	0	0
L_p -loss, $p \in (1, 2)$	$ t ^p$	$\frac{p(p-1)}{8} B^{p-2} \varepsilon^2$	$\frac{p}{8(p-1)^2} B^{p-2} \varepsilon^2$
L_p -loss, $p \in [2, \infty)$	$ t ^p$	$(\frac{\varepsilon}{2})^p$	$(\frac{\varepsilon}{2})^p$
Logistic loss	$-\ln \frac{4e^t}{(1+e^t)^2}$	$\frac{e^{\varepsilon/2}-1}{2e^{\varepsilon/2}} \ln \frac{e^B+e^\varepsilon}{e^B+e^{\varepsilon/2}}$	$\frac{e^{\varepsilon/2}-1}{e^{\varepsilon/2}} \ln \frac{e^B+e^\varepsilon}{e^B+e^{\varepsilon/2}}$
Huber's loss, $c > 0$	$\begin{cases} \frac{t^2}{2} & \text{if } t \leq c \\ c t - \frac{c^2}{2} & \text{else} \end{cases}$	$\begin{cases} \frac{\varepsilon^2}{8} & \text{if } B \leq c \\ 0 & \text{else} \end{cases}$	$\begin{cases} \frac{\varepsilon^2}{8} & \text{if } B \leq c \\ 0 & \text{else} \end{cases}$

Table 3: Some invariant loss functions and corresponding upper and lower bounds of $\delta_{\psi|[-B,B]}(\varepsilon)$, $0 < \varepsilon \leq B$, for the restriction $\psi|_{[-B,B]}$ of ψ to $[-B, B]$, for $B > 0$. The asymptotic behaviour for the L_p -loss, $1 < p < 2$, is computed in Example A.4. For the L_p -loss, $p \geq 2$, and Huber's loss the lower bounds can be found by Clarkson's inequality, and the upper bounds were found by finding suitable $x_1, x_2 \in [B, B]$. The calculations for the logistic loss can be found in Example A.5.

In Theorem 3.21 we have seen that for $Q \in \mathcal{Q}_{\text{sym}}^*(L)$ we may have $\delta_{\max}(\varepsilon, Q) > 0$, $\varepsilon > 0$, even if L is not strictly convex. The key reason for this possibility was the assumption that Q has some mass around its center. Now recall that in the proof of the upper bounds of Theorem 3.23 we used the fact that for general $Q \in \mathcal{Q}_{\text{sym}}^*$ this mass can be arbitrarily small. However, if we enforce lower bounds on this mass the construction of, this proof no longer works. Instead, it turns out that we can establish lower bounds on $\delta_{\max}(\varepsilon, Q)$ as the following example illustrate:

Example 3.25 Let L be the distance-based loss function whose associated ψ is $\psi(t) = |t|$, $t \in \mathbb{R}$. Then for all $Q \in \mathcal{Q}_{\text{sym}}(L)$, $m := \mathbb{E}Q$, and all $\varepsilon > 0$ we have

$$\delta_{\max}(\varepsilon, Q) = \int_0^\varepsilon Q^{(m)}((-s, s)) ds.$$

To see this we first observe that for $t \geq 0$ and $y \in \mathbb{R}$ an easy calculation shows

$$\psi(y-t) - \psi(y) = \begin{cases} t & \text{if } y \leq 0 \\ t-2y & \text{if } y \in (0, t) \\ -t & \text{if } y \geq t. \end{cases}$$

Consequently the symmetry of $Q^{(m)}$ yields

$$\begin{aligned} \mathcal{C}_{L,Q}(m+t) - \mathcal{C}_{L,Q}^* &= tQ^{(m)}((-\infty, 0]) + tQ^{(m)}((0, t)) - 2 \int_{0 < y < t} y dQ^{(m)}(y) - tQ^{(m)}([t, \infty)) \\ &= tQ^{(m)}((-t, t)) - 2 \int_{0 \leq y < t} y dQ^{(m)}(y) \\ &= \int_0^t Q^{(m)}((-t, t)) - Q^{(m)}((-t, -s] \cup [s, t)) ds \\ &= \int_0^t Q^{(m)}((-s, s)) ds. \end{aligned}$$

From this we easily find the assertion by Lemma 3.19.

Remark 3.26 The above results show that using surrogate loss functions for regression problems requires some care: for example let us suppose that the primary goal of the regression problem is to find the regression function. If we only know that the noise distributions have finite variances (and expect that these distributions are rather asymmetric) then the least squares loss is the only reasonable choice by Proposition 3.17. However, if we know that the noise is (almost) symmetric then e.g. symmetric, strictly convex and Lipschitz continuous losses like the logistic loss can be a reasonable alternative. In addition, if we are confident that the noise is rather concentrated around its mean, e.g. in the form of $Q^{(m)}((-s, s)) > c_Q s^q$ for small $s > 0$, then even convex loss functions like the absolute distance loss considered in the previous example can be a good choice. Finally, if we additionally expect that the data set contains extreme outliers then the logistic loss or the absolute distance loss may actually be the better choice than the least squares loss. However, recall that such a decision only makes sense under an (almost) symmetric behaviour of the noise distribution.

When we introduced the template loss L_{mean} we also discussed its relation to self-calibration issues. Therefore let us finally investigate in which sense convex, distance-based loss functions are self-calibrated.

Theorem 3.27 *Let L be a distance-based, convex loss function whose associated ψ satisfies $\lim_{t \rightarrow \pm\infty} \psi(t) = \infty$. Then we have:*

- i) L is self-calibrated with respect to $\mathcal{Q}_{\text{bounded}}$.
- ii) If L is of some growth type $p \geq 1$ then L is self-calibrated with respect to $\mathcal{Q}_{\mathbb{R}}(L)$.
- iii) If L is strictly convex then L is uniformly self-calibrated with respect to $\mathcal{Q}_{[-M, M]}$ for all $M > 0$, and we have

$$\delta_{\max, \check{L}, L}(\varepsilon, \mathcal{Q}_{[-M, M]}) \geq 2\delta_{\psi|_{[-(2M+\varepsilon), 2M+\varepsilon]}}(\varepsilon), \quad \varepsilon > 0, M > 0. \quad (45)$$

- iv) If L is uniformly convex then L is uniformly self-calibrated with respect to $\mathcal{Q}_{\text{bounded}}$ and we have

$$\delta_{\max, \check{L}, L}(\varepsilon, \mathcal{Q}_{\text{bounded}}) \geq 2\delta_{\psi}(\varepsilon), \quad \varepsilon > 0. \quad (46)$$

- v) If L is uniformly convex and of some growth type $p \geq 2$ then L is uniformly self-calibrated with respect to $\mathcal{Q}_{\mathbb{R}}(L)$ and we have

$$\delta_{\max, \check{L}, L}(\varepsilon, \mathcal{Q}_{\mathbb{R}}(L)) \geq 2\delta_{\psi}(\varepsilon), \quad \varepsilon > 0. \quad (47)$$

Proof: i). Proposition 3.16 shows $\mathcal{Q}_{\text{bounded}} \subset \mathcal{Q}_{\min}(L)$, and hence Lemma 2.36 implies i).
ii). The results in [8] together with Proposition 3.16 show $\mathcal{Q}_{\mathbb{R}}(L) \subset \mathcal{Q}_{\min}(L)$. Again, the assertion then follows from Lemma 2.36.
iii). Let us fix a $Q \in \mathcal{Q}_{[-M, M]}$. By Proposition 3.16 we then know that there exists a $t_Q^* \in \mathbb{R}$ with $\mathcal{M}_{L, Q}(0^+) = \{t_Q^*\}$, and consequently (8) reduces to

$$\delta_{\max, \check{L}, L}(\varepsilon, Q) = \min\{\mathcal{C}_{L, Q}(t_Q^* - \varepsilon) - \mathcal{C}_{L, Q}^*, \mathcal{C}_{L, Q}(t_Q^* + \varepsilon) - \mathcal{C}_{L, Q}^*\}. \quad (48)$$

Let us now assume that $t_Q^* > M$. Then we have $y - t_Q^* \leq y - M \leq 0$ for all $y \in [-M, M]$. Since ψ is convex and has a minimum at 0 it is decreasing on $(-\infty, 0]$ and hence we find $\psi(y - M) \leq \psi(y - t_Q^*)$ for all $y \in [-M, M]$. This directly implies $\mathcal{C}_{L, Q}(M) \leq \mathcal{C}_{L, Q}(t_Q^*)$, and hence our assumption $t_Q^* > M$ cannot be true. Since we can analogously show $t_Q^* \geq -M$, we have $t_Q^* \in [-M, M]$. Moreover, for

$t \in [-M - \varepsilon, M + \varepsilon]$ we have

$$\begin{aligned}
\frac{\mathcal{C}_{L,Q}(t) + \mathcal{C}_{L,Q}^*}{2} &= \int_{[-M,M]} \frac{\psi(y-t) + \psi(y-t_Q^*)}{2} dQ(y) \\
&\geq \int_{[-M,M]} \psi\left(y - \frac{t+t_Q^*}{2}\right) + \delta_{\psi|_{[-(2M+\varepsilon), 2M+\varepsilon]}}(|t-t_Q^*|) dQ(y) \\
&= \mathcal{C}_{L,Q}\left(\frac{t+t_Q^*}{2}\right) + \delta_{\psi|_{[-(2M+\varepsilon), 2M+\varepsilon]}}(|t-t_Q^*|) \\
&\geq \mathcal{C}_{L,Q}^* + \delta_{\psi|_{[-(2M+\varepsilon), 2M+\varepsilon]}}(|t-t_Q^*|), \tag{49}
\end{aligned}$$

and thus we find (45) by combining (48) with (49) and $t_Q^* \in [-M, M]$.

iv). Inequality (45) immediately implies (46).

v). The proof of Inequality (47) is analogous to the proof of (45). ■

Remark 3.28 The above theorem shows that for convex, distance-based losses L , approximate L -risk minimizers approximate the Bayes decision functions in the sense of Theorem 2.37, i.e. in probability. In particular note that the absolute distance loss can be used to estimate the median (multi)-function in this weak sense. Moreover, Example 3.25 together with Theorem 2.18 shows that in order to estimate the median function in a stronger sense one needs assumptions on the concentration of the noise distributions. The deeper reason for this observation is the fact that for *symmetric, convex* losses L whose associated ψ have a *unique* minimum at 0 the stronger convexity assumptions in *iii-v*) are also *necessary* for uniform self-calibration. To see the latter observe that for such L Theorem 3.21 shows the L_{mean} -calibration with respect to $\mathcal{Q}_{\text{sym}}^*(L)$, and from Lemma 3.19 we may thus conclude that $\mathcal{C}_{L,Q}(\cdot)$ has a *unique* minimum at $\mathbb{E}Q$ for all $Q \in \mathcal{Q}_{\text{sym}}^*(L)$. Equation (39) then shows

$$\delta_{\max, L_{\text{mean}}, L}(\varepsilon, Q) = \delta_{\max, \check{L}, L}(\varepsilon, Q), \quad \varepsilon \in [0, \infty], Q \in \mathcal{Q}_{\text{sym}}^*(L). \tag{50}$$

Now assume e.g. that L is uniformly self-calibrated with respect to $\mathcal{Q}_{[-M,M]}$ for all $M > 0$. Then it is also uniformly self-calibrated with respect to $\mathcal{Q}_{[-M,M], \text{sym}}^*$, and thus (50) shows that L is L_{mean} -calibrated with respect to $\mathcal{Q}_{[-M,M], \text{sym}}^*$. By Theorem 3.23 we then see that L is strictly convex.

Remark 3.29 Inequalities for excess risks resulting from Theorem 3.27 can be used to establish *variance bounds* which are important for bounding the estimation error by Talagrand's inequality. Indeed, if e.g. L is a strictly convex loss function then its corresponding ψ is locally Lipschitz continuous and hence (28) can be used to find an “inner” version of a variance bound. This approach was somewhat implicitly taken in e.g. [1, 4] in order to derive variance bounds for margin-based losses, but of course it also works for distance-based losses. However, it sometimes provides too loose bounds as e.g. the hinge loss shows. Indeed, this loss fails to be uniformly self-calibrated not only for $\eta \rightarrow 1/2$ but also for $\eta \rightarrow 0$ and $\eta \rightarrow 1$. In order to establish variance bounds using the self-calibration we consequently *need* not only to ensure a noise assumption in the sense of Tsybakov, but also a similar assumption ensuring that the set where η is either close to 0 or 1 is small. However, the latter is superfluous as the variance bound established in [28] shows. One may ask whether variance bounds can also be *directly* established using our general framework and a (template) loss which reflects the variance bounds one is interested in. Some preliminary considerations we have already made in this directions are promising but due to space constraints we do not go into further details.

3.4 Density Level Detection

In this subsection we show how the developed theory can be used to investigate the density level detection problem. To this end let us first recall that in this *unsupervised* learning problem we have a *known* probability measure μ on X and the goal is to estimate a level set $\{g > \rho\}$ or $\{g \geq \rho\}$ of an *unknown* function $g : X \rightarrow [0, \infty)$ satisfying $\|g\|_{L_1(\mu)} = 1$. The only information to achieve this goal is given to us by a data set $T := (x_1, \dots, x_n)$ of n samples drawn independently from the probability measure $g\mu$. Typically, an estimate of a level set is of the form $\{f > 0\}$, where

$f : X \rightarrow \mathbb{R}$ is a measurable function. In order to assess the quality of such an estimate one can use the loss function $L_{\text{DLD}} : X \times \mathbb{R} \rightarrow [0, \infty]$ defined by

$$L_{\text{DLD}}(x, t) := \mathbf{1}_{(-\infty, 0)}((g(x) - \rho) \text{sign } t), \quad x \in X, t \in \mathbb{R}. \quad (51)$$

Note that this loss function penalizes predictions t if either $t \geq 0$ and $g(x) < \rho$, or $t < 0$ and $g(x) > \rho$, whereas it completely ignores t if $g(x) = \rho$. In a slight abuse of notations we now write

$$\mathcal{R}_{L_{\text{DLD}}, \mu}(f) := \int_X L_{\text{DLD}}(x, f(x)) d\mu(x),$$

where $f : X \rightarrow \mathbb{R}$ is a measurable function. Obviously, this definition gives $\mathcal{R}_{L_{\text{DLD}}, \mu}(f) = \mathcal{R}_{L_{\text{DLD}}, P}(f)$ for every Polish space Y and every distribution P on $X \times Y$ with $P_X = \mu$. Moreover note, that if we additionally assume $\mu(\{g = \rho\}) = 0$, as it is usually done in the literature³, we obtain the better known expression

$$\mathcal{R}_{L_{\text{DLD}}, \mu}(f) = \mu(\{g \geq \rho\} \triangle \{f \geq 0\}).$$

Now note that unlike for the supervised loss functions we *cannot compute* $L_{\text{DLD}}(x, t)$, since g is unknown to us, and consequently for training sets of the form $T = (x_1, \dots, x_n) \in X^n$ we cannot use e.g. an ERM based on L_{DLD} simply because we cannot compute the empirical risk. To overcome this problem a framework was developed in [27] that translates the density level detection problem to a binary classification problem. The key idea of this framework was the following definition:

Definition 3.30 *Let μ be a distribution on X and $Y := \{-1, 1\}$. Furthermore, let $g : X \rightarrow [0, \infty]$ be a measurable function with $\|g\|_{L_1(\mu)} = 1$. For $\rho > 0$ we define the distribution $P := g\mu \ominus_\rho \mu$ by*

$$\begin{aligned} P_X &:= \frac{g + \rho}{1 + \rho} \mu, \\ P(y = 1|x) &:= \frac{g(x)}{g(x) + \rho}, \quad x \in X. \end{aligned}$$

Note that the distribution $P := g\mu \ominus_\rho \mu$ describes a binary classification problem and consequently one could ask in which way the classification risk $\mathcal{R}_{L_{\text{class}}, P}(\cdot)$ is related to $\mathcal{R}_{L_{\text{DLD}}, \mu}(\cdot)$. It turned out in [27] that there is indeed a close relationship in between the two quantities. The goal in this subsection is to extend the considerations of [27] using our general theory. To this end we write $Y := \{-1, 1\}$ throughout this section. Furthermore, we always identify a distribution $Q \in \mathcal{Q}_Y$ by $\eta \in [0, 1]$ via the relationship $\eta := Q(\{1\})$. Now let us define $\bar{L}_{\text{DLD}} : [0, 1] \times \mathbb{R} \rightarrow [0, \infty]$ by

$$\bar{L}_{\text{DLD}}(\eta, t) := (1 - \eta) \mathbf{1}_{(-\infty, 0)}((2\eta - 1) \text{sign } t). \quad (52)$$

Using the identification $\eta = Q(\{1\})$ we can then think of \bar{L}_{DLD} as a *template* loss. Furthermore, for $P = g\mu \ominus_\rho \mu$ the P -instance $\bar{L}_{\text{DLD}, P}$ of \bar{L}_{DLD} is

$$\bar{L}_{\text{DLD}, P}(x, t) = (1 - \eta(x)) \mathbf{1}_{(-\infty, 0)}((2\eta(x) - 1) \text{sign } t) = \frac{\rho}{g(x) + \rho} \mathbf{1}_{(-\infty, 0)}((g(x) - \rho) \text{sign } t),$$

where $\eta(x) := P(y = 1|x) = \frac{g(x)}{g(x) + \rho}$. With this equation we then find

$$\mathcal{R}_{L_{\text{DLD}}, \mu}(f) = \int_X L_{\text{DLD}}(x, f(x)) \frac{1 + \rho}{g(x) + \rho} dP_X(x) = \frac{1 + \rho}{\rho} \mathcal{R}_{\bar{L}_{\text{DLD}, P}, P}(f)$$

³We could have gone this path, too. However, we will see that technically it is a bit more convenient to ignore the set $\{g = \rho\}$ rather than to assume that it (essentially) does not exist.

for all measurable $f : X \rightarrow \mathbb{R}$. Consequently, suitable supervised surrogates for the DLD-problem are exactly the losses that are (uniformly) \bar{L}_{DLD} -calibrated. In order to identify the latter losses let us first compute the corresponding calibration functions:

Lemma 3.31 *Let $L : Y \times \mathbb{R} \rightarrow [0, \infty]$ be a supervised loss function. Then for all $\eta \in [0, 1]$ and $\varepsilon \in (0, \infty]$ we have*

$$\delta_{\max, \bar{L}_{\text{DLD}}, L}(\varepsilon, \eta) = \begin{cases} \infty & \text{if } \varepsilon > 1 - \eta \\ \inf_{t \in \mathbb{R}: (2\eta-1) \operatorname{sign} t < 0} \mathcal{C}_{L, \eta}(t) - \mathcal{C}_{L, \eta}^* & \text{if } \varepsilon \leq 1 - \eta. \end{cases}$$

Proof: A simple calculation shows $\mathcal{C}_{\bar{L}_{\text{DLD}}, \eta}^* = 0$, and consequently we obtain $\mathcal{M}_{\bar{L}_{\text{DLD}}, \eta}(\varepsilon) = \mathbb{R}$ if $\varepsilon > 1 - \eta$, and $\mathcal{M}_{\bar{L}_{\text{DLD}}, \eta}(\varepsilon) = \{t \in \mathbb{R} : (2\eta - 1) \operatorname{sign} t \geq 0\}$ otherwise. ■

With the help of the above lemma we now obtain the first main result which compares classification calibration with \bar{L}_{DLD} -calibration:

Theorem 3.32 *Let $L : Y \times \mathbb{R} \rightarrow [0, \infty]$ be a supervised loss function, $\eta \in [0, 1]$, and $0 \leq \varepsilon \leq \min\{1 - \eta, |2\eta - 1|\}$. Then we have*

$$\delta_{\max, \bar{L}_{\text{DLD}}, L}(\varepsilon, \eta) \geq \delta_{\max, L_{\text{class}}, L}(\varepsilon, \eta),$$

and consequently, L is \bar{L}_{DLD} -calibrated if L is classification calibrated. Moreover, if L is continuous then the above inequality becomes an equality and L is classification calibrated if and only if L is \bar{L}_{DLD} -calibrated.

Proof: Combining Lemma 3.1 with Lemma 3.31 yields

$$\delta_{\max, \bar{L}_{\text{DLD}}, L}(\varepsilon, \eta) = \inf_{\substack{t \in \mathbb{R}: \\ (2\eta-1) \operatorname{sign} t < 0}} \mathcal{C}_{L, \eta}(t) - \mathcal{C}_{L, \eta}^* \geq \inf_{\substack{t \in \mathbb{R}: \\ (2\eta-1) \operatorname{sign} t \leq 0}} \mathcal{C}_{L, \eta}(t) - \mathcal{C}_{L, \eta}^* = \delta_{\max, L_{\text{class}}, L}(\varepsilon, \eta).$$

Moreover, for continuous L the assertion can be found using the continuity of $t \mapsto \mathcal{C}_{L, \eta}(t)$. ■

By the results on classification calibrated, margin-based loss functions in [1] we immediately obtain a variety of \bar{L}_{DLD} -calibrated losses. Furthermore, the P -instances of \bar{L}_{DLD} are bounded loss functions and therefore Theorem 2.8 yields

$$\mathcal{R}_{L, P}(f_n) \rightarrow \mathcal{R}_{L, P}^* \implies \mathcal{R}_{L_{\text{DLD}}, \mu}(f_n) \rightarrow 0$$

whenever $P = g\mu \ominus_{\rho} \mu$ and L is classification calibrated. In addition, for $L := L_{\text{class}}$ the proof of [27, Thm. 4] shows that the converse implication is also true.

Our next goal is to identify the *uniformly* \bar{L}_{DLD} -calibrated losses. The following theorem gives a complete, though rather disappointing solution:

Theorem 3.33 *There exists no supervised loss function $L : Y \times \mathbb{R} \rightarrow [0, \infty]$ that is uniformly \bar{L}_{DLD} -calibrated.*

Proof: Let $L : Y \times \mathbb{R} \rightarrow [0, \infty]$ be a supervised loss function, and $\varepsilon := \frac{1}{4}$. For $\eta \in [0, \frac{1}{2}]$ we define

$$H(\eta) := \inf_{t \in \mathbb{R}: (2\eta-1)t \leq 0} \mathcal{C}_{L, \eta}(t) - \mathcal{C}_{L, \eta}^*.$$

Then we obviously have $H(1/2) = 0$, and Lemma 3.31 together with $\text{sign } 0 = 1$ yields

$$\delta_{\max, \bar{L}_{\text{DLD}}, L}(1/4, \eta) = H(\eta), \quad 0 \leq \eta < 1/2.$$

Furthermore, both $\eta \mapsto \mathcal{C}_{L, \eta}^*$ and $\eta \mapsto \inf_{t \in \mathbb{R}: (2\eta-1)t \leq 0} \mathcal{C}_{L, \eta}(t)$ are defined by infimums taken over affine functions of η and thus they are concave and upper semi-continuous. Since in addition they are both finite on $[0, 1/2]$, [23, Cor. 2.37] shows that H is continuous. Consequently, we obtain

$$\delta_{\max, \bar{L}_{\text{DLD}}, L}(1/4, \mathcal{Q}_Y) \leq \inf_{\eta \in [0, 1/2)} \delta_{\max, \bar{L}_{\text{DLD}}, L}(1/4, \eta) = \inf_{\eta \in [0, 1/2)} H(\eta) = 0.$$

■

Remark 3.34 Theorem 2.18 showed that uniform calibration is necessary to establish inequalities between excess risks if essentially no assumptions on the data-generating distribution are imposed.⁴ Together with Theorem 3.33 we consequently see that it is *impossible* to find a supervised loss $L : Y \times \mathbb{R} \rightarrow [0, \infty]$ and an increasing function $\delta : [0, \infty] \rightarrow [0, \infty]$ with $\delta(0) = 0$ and $\delta(\varepsilon) > 0$ for $\varepsilon > 0$, such that

$$\delta(\mathcal{R}_{L_{\text{DLD}}, \mu}(f)) \leq \mathcal{R}_{L, P}(f) - \mathcal{R}_{L, P}^*$$

holds for all possible choices of f, g, ρ and $P := g\mu \ominus_\rho \mu$. Furthermore, the proof of Theorem 3.33 shows that this observation remains true even if we only consider *one fixed* level ρ and additionally exclude pairs (μ, g) with $\mu(\{g = \rho\}) = 0$. On the other hand if we only consider densities g that are bounded away from the level ρ then the corresponding function $x \mapsto \eta(x)$ is bounded away from the critical level $\frac{1}{2}$, and thus the argument of Theorem 3.33 does not work. This observation has been implicitly used for the inequalities established in [27, Thm. 10]. Moreover, using Theorem 2.30 we can actually improve these inequalities slightly. However, due to space constraints we omit the details.

3.5 Density Estimation

In this last subsection we apply the developed theory to the density estimation problem. To this end let us first recall that in this *unsupervised* learning problem we have a data-generating distribution on X which is of the form $g\mu$, where μ is a known distribution and g is an unknown density. The learning goal is then to estimate g . Let us assume that we have an estimate f of this density. Then the quality of this estimate is usually measured by

$$\mathcal{R}_{L_g, \mu}(f) := \int_X |f(x) - g(x)| d\mu(x) = \|f - g\|_{L_1(\mu)},$$

or if $g \in L_p(\mu)$ is known, by $\mathcal{R}_{\bar{L}_{g, p}, \mu}(f) := \|f - g\|_{L_p(\mu)}$. Obviously, the above performance measure $\mathcal{R}_{L_g, \mu}(\cdot)$ is a risk with respect to the unsupervised loss function $(x, t) \mapsto |t - g(x)|$. However, like for the density level detection problem, this loss function is not accessible to us since we do not know g . Consequently, if we want to assess the quality of an estimate, we need a surrogate risk, i.e. a surrogate loss function. In order to find such a surrogate recall that a well-known heuristic for the density estimation problem is based on using additional samples drawn from μ (see e.g. [12, Chap. 14.2.4]). Let us now briefly describe and analyze this approach. To this we define $P := g\mu \ominus_1 \mu$ on $X \times Y$, $Y := \{-1, 1\}$, i.e. P is the joint distribution of the original, positively labeled samples drawn from $g\mu$, and the artificial, negatively labeled samples drawn from μ . As usual, we identify the conditional probability $P(\cdot|x)$ with $\eta(x)$, via the relation $P(y = 1|x) = \eta(x)$, so that we have

$$\eta(x) = \frac{g(x)}{1 + g(x)} \quad \text{and} \quad g(x) = \frac{\eta(x)}{1 - \eta(x)}.$$

⁴Formally the result only holds for loss functions, *not template losses*. However, it is quite straightforward to see that the proof of Theorem 2.18 can be easily modified to establish an analogous result for instances of template losses.

Moreover, let $L : Y \times \overline{\mathbb{R}} \rightarrow [0, \infty)$ be a regular loss function such that $\mathcal{M}_{L,\eta}(0^+)$ contains a single element, denoted by $m(\eta)$, for every $\eta \in [0, 1)$. Moreover, assume that the map $m : [0, 1) \rightarrow M := \{m(\eta) : \eta \in [0, 1)\}$ defined by $\eta \mapsto m(\eta)$ is a bijection, and that M is a Polish space⁵. Given a $t \in M$ one can then interpret $m^{-1}(t)$ as an estimate of η , and therefore the template loss

$$\tilde{L}(\eta, t) := 2(1 - \eta) \left| \frac{m^{-1}(t)}{1 - m^{-1}(t)} - \frac{\eta}{1 - \eta} \right|$$

measures the quality of the estimate $\frac{m^{-1}(t)}{1 - m^{-1}(t)}$ for g . Moreover, for measurable $f : X \rightarrow M$ an easy calculation shows

$$\mathcal{R}_{\tilde{L},P}(f) = \int_X 2(1 - \eta(x)) \left| \frac{m^{-1}(f(x))}{1 - m^{-1}(f(x))} - \frac{\eta(x)}{1 - \eta(x)} \right| dP_X(x) = \int_X \left| \frac{m^{-1}(f(x))}{1 - m^{-1}(f(x))} - g(x) \right| d\mu(x),$$

i.e. we have $\mathcal{R}_{\tilde{L},P}(f) = \mathcal{R}_{L_g,\mu}(\frac{m^{-1}(f)}{1 - m^{-1}(f)})$. Consequently, it suffices to investigate surrogates for the template loss \tilde{L} . Let us begin with a negative result:

Theorem 3.35 *There exists no regular loss function satisfying the above additional assumptions, such that L is uniformly \tilde{L} -calibrated.*

Proof: Let us fix $\eta \in [0, 1)$ and $\varepsilon > 0$. Then we have $\frac{2\eta + \varepsilon}{2 + \varepsilon} \in [0, 1)$ and consequently $t_\eta := m(\frac{2\eta + \varepsilon}{2 + \varepsilon})$ is well-defined. In addition, an easy calculation shows $\tilde{L}(\eta, t_\eta) = \varepsilon$ and hence we obtain

$$\delta_{\max, \tilde{L}, L}(\varepsilon, \eta) = \inf_{t \in M : \tilde{L}(\eta, t) \geq \varepsilon} \mathcal{C}_{L,\eta}(t) - \mathcal{C}_{L,\eta}^* \leq \mathcal{C}_{L,\eta}(t_\eta) - \mathcal{C}_{L,\eta}^*.$$

Moreover, we have

$$\begin{aligned} \mathcal{C}_{L,\eta}(t_\eta) &= \eta L(1, t_\eta) + (1 - \eta) L(-1, t_\eta) \\ &= \frac{2\eta + \varepsilon}{2 + \varepsilon} L(1, t_\eta) + \left(1 - \frac{2\eta + \varepsilon}{2 + \varepsilon}\right) L(-1, t_\eta) + \frac{\varepsilon(\eta - 1)}{2 + \varepsilon} L(1, t_\eta) + \frac{\varepsilon(1 - \eta)}{2 + \varepsilon} L(-1, t_\eta) \\ &= \mathcal{C}_{L, \frac{2\eta + \varepsilon}{2 + \varepsilon}}^* + \frac{\varepsilon}{2 + \varepsilon} \mathcal{C}_{L,\eta}(t_\eta) - \frac{\varepsilon}{2 + \varepsilon} \mathcal{C}_{L,1}(t_\eta), \end{aligned}$$

and consequently we obtain

$$\mathcal{C}_{L,\eta}(t_\eta) = \frac{2 + \varepsilon}{2} \mathcal{C}_{L, \frac{2\eta + \varepsilon}{2 + \varepsilon}}^* - \frac{\varepsilon}{2} \mathcal{C}_{L,1}(t_\eta) \leq \frac{2 + \varepsilon}{2} \mathcal{C}_{L, \frac{2\eta + \varepsilon}{2 + \varepsilon}}^* - \frac{\varepsilon}{2} \mathcal{C}_{L,1}^*.$$

Together with our estimate on the calibration function this yields

$$\delta_{\max, \tilde{L}, L}(\varepsilon, \eta) \leq \frac{2 + \varepsilon}{2} \mathcal{C}_{L, \frac{2\eta + \varepsilon}{2 + \varepsilon}}^* - \frac{\varepsilon}{2} \mathcal{C}_{L,1}^* - \mathcal{C}_{L,\eta}^*.$$

Now recall that we have already seen in the proof of Theorem 3.33, that $\eta \mapsto \mathcal{C}_{L,\eta}^*$ is continuous on $[0, 1]$ and hence we find $\lim_{\eta \rightarrow 1} \delta_{\max, \tilde{L}, L}(\varepsilon, \eta) = 0$. From this we easily infer the assertion. \blacksquare

Remark 3.36 At first glance one is tempted to conclude from Theorem 3.35 and Theorem 2.18 that there cannot exist a general inequality between the risk $\mathcal{R}_{L_g,\mu}(\frac{m^{-1}(f)}{1 - m^{-1}(f)})$ and an excess surrogate risk $\mathcal{R}_{L,P}(f) - \mathcal{R}_{L,P}^*$. However, in the case of density estimation we are actually not interested in all distributions P on $X \times Y$, but only in distributions of the form $g\mu \ominus_1 \mu$ with g being a *density* with respect to μ , and therefore

⁵This is only a technical assumption which is satisfied for all commonly used surrogates.

Theorem 2.18 does not apply *directly*. Nevertheless, Theorem 3.35 can be used to show that no supervised loss L satisfying the additional assumptions on $m(\cdot)$ allow a general inequality between $\mathcal{R}_{L_g, \mu}(\frac{m^{-1}(f)}{1-m^{-1}(f)})$ and its excess risk $\mathcal{R}_{L,P}(f) - \mathcal{R}_{L,P}^*$. To see let $X := \{1, 2\}$, μ be a distribution on X and $g : X \rightarrow [0, \infty)$ be a density with $g(2) = 0$. Note that μ is uniquely determined by $\mu_1 := \mu(\{1\})$, and since g is a density with respect to μ we have $g_1 := g(1) = 1/\mu_1$. Using our standard notations this yields $\eta_1 := \eta(1) \geq 1/2$ and $\eta(2) = 0$. Now note that for $f : X \rightarrow \mathbb{R}$ with $m^{-1}(f(2)) = 0 = g(2)$ we have

$$\mathcal{R}_{L,P}(f) - \mathcal{R}_{L,P}^* = \frac{\mathcal{C}_{L, \eta_1}(f_1) - \mathcal{C}_{L, \eta_1}^*}{2\eta_1} \leq \mathcal{C}_{L, \eta_1}(f_1) - \mathcal{C}_{L, \eta_1}^*,$$

where $f_1 := f(1)$. Moreover, for such f we also have

$$\mathcal{R}_{L_g, \mu}\left(\frac{m^{-1}(f)}{1-m^{-1}(f)}\right) = \frac{1-\eta_1}{\eta_1} \cdot \left| \frac{m^{-1}(f_1)}{1-m^{-1}(f_1)} - \frac{\eta_1}{1-\eta_1} \right| = \frac{\tilde{L}(\eta_1, f_1)}{2\eta_1}.$$

Now assume that we have a function $\delta : [0, \infty) \rightarrow [0, \infty]$ and an inequality in the sense of Theorem 2.18. Then this inequality in particular holds for $f : X \rightarrow \mathbb{R}$ with $m^{-1}(f(2)) = 0$, and hence we find

$$\delta\left(\frac{\tilde{L}(\eta_1, f_1)}{2}\right) \leq \delta\left(\frac{\tilde{L}(\eta_1, f_1)}{2\eta_1}\right) = \delta\left(\mathcal{R}_{L_g, \mu}\left(\frac{m^{-1}(f)}{1-m^{-1}(f)}\right)\right) \leq \mathcal{R}_{L,P}(f) - \mathcal{R}_{L,P}^* \leq \mathcal{C}_{L, \eta_1}(f_1) - \mathcal{C}_{L, \eta_1}^*.$$

Since this inequality holds for all $\eta_1 \in [1/2, 1)$ and all $f_1 \in \mathbb{R}$, we easily see that L is uniformly \tilde{L} -calibrated with respect to $\{\eta : 1/2 \leq \eta < 1\}$. However, we have seen in (the proof of) Theorem 3.35 that this uniform calibration is impossible, and consequently no general inequalities are possible in the above sense.

After these disappointing results let us finally present a positive result for convex loss functions:

Theorem 3.37 *Let $L : Y \times \mathbb{R} \rightarrow [0, \infty)$ be a regular loss function which is convex on in its second argument. If L satisfies the above additional assumptions then L is \tilde{L} -calibrated.*

Proof: Repeating the proof of [25, Lem. 20] we see that the map $m : [0, 1) \rightarrow M$ is monotone, and since it is also invertible, it must be strictly monotone. Moreover, we obviously have $\mathcal{M}_{L, Q, x}(0^+) \subset \mathcal{M}_{\tilde{L}, Q, x}(0^+)$. Using Lemma 2.11 and the ideas of its proof we then obtain the assertion. ■

Remark 3.38 Let $L : Y \times \mathbb{R} \rightarrow [0, \infty)$ be regular loss function that is \tilde{L} -calibrated. Repeating the proof of Theorem 2.37 we then see that

$$\frac{m^{-1}(f_n)}{1-m^{-1}(f_n)} \rightarrow g \quad \text{in probability}$$

whenever $f_n : X \rightarrow M$ is a sequence of measurable functions with $\mathcal{R}_{L,P}(f_n) \rightarrow \mathcal{R}_{L,P}^*$. In particular, every universally L -risk consistent method is also universally consistent in the above weak form of the density estimation problem. Moreover, if g is a bounded density, then it is not hard to derive consistency with respect to $\mathcal{R}_{L_g, \mu}(\cdot)$. Finally, we like to mention that additional results can be obtained by computing the calibration function $\delta_{\max, \tilde{L}, L}(\cdot, \cdot)$ for specific L , but because of space constraints we omit the details.

4 Proofs of the Results of the General Theory

Proof of Lemma 2.5: For $n \in \mathbb{N}$ let us fix an $\alpha_n \in \mathcal{A}$ that satisfies (2) for $\varepsilon := 1/n$. Then for all $x \in X$ we have

$$\mathcal{C}_{L, P(\cdot|x), x}^* = \inf_{\alpha \in \mathcal{A}} \mathcal{C}_{L, P(\cdot|x), x}(\alpha) = \inf_{n \in \mathbb{N}} \mathcal{C}_{L, P(\cdot|x), x}(\alpha_n),$$

and hence $x \mapsto \mathcal{C}_{L, P(\cdot|x), x}^*$ is measurable. Furthermore for all $n \geq 1$ we have

$$\begin{aligned} \mathcal{R}_{L,P}^* &\leq \int_X \mathcal{C}_{L, P(\cdot|x), x}(\alpha_n) dP_X(x) \leq \int_X \mathcal{C}_{L, P(\cdot|x), x}^* dP_X(x) + \frac{1}{n} \leq \inf_{\alpha \in \mathcal{A}} \int_X \mathcal{C}_{L, P(\cdot|x), x}(\alpha) dP_X(x) + \frac{1}{n} \\ &\leq \mathcal{R}_{L,P}^* + \frac{1}{n}, \end{aligned}$$

and hence we obtain the second assertion if we let $n \rightarrow \infty$. \blacksquare

Proof of Theorem 2.8: For brevity's sake we write $\mathcal{C}_{1,x}(\alpha) := \mathcal{C}_{L_1,P(\cdot|x),x}(\alpha) - \mathcal{C}_{L_1,P(\cdot|x),x}^*$ and $\mathcal{C}_{2,x}(\alpha) := \mathcal{C}_{L_2,P(\cdot|x),x}(\alpha) - \mathcal{C}_{L_2,P(\cdot|x),x}^*$ for $x \in X$, $\alpha \in \mathcal{A}$. Furthermore let us fix an $\varepsilon > 0$ and write $h(x) := \delta(\varepsilon, x)$, $x \in X$. Then for all $x \in X$, and all $\alpha \in \mathcal{A}$ with $\mathcal{C}_{1,x}(\alpha) \geq \varepsilon$ we have $\mathcal{C}_{2,x}(\alpha) \geq h(x)$, and hence we obtain

$$\mathcal{R}_{L_2,P}(\alpha) - \mathcal{R}_{L_2,P}^* = \int_X \mathcal{C}_{2,x}(\alpha) dP_X(x) \geq \int_{\mathcal{C}_{1,x}(\alpha) \geq \varepsilon} h(x) dP_X(x)$$

for all $\alpha \in \mathcal{A}$. Furthermore, since $h(x) > 0$ holds for all $x \in X$, the measure $\nu := b P_X$ is absolutely continuous with respect to $\mu := h P_X$ and consequently there exists a $\delta > 0$ such that $\nu(A) \leq \varepsilon$ for all measurable $A \subset X$ with $\mu(A) \leq \delta$. For $\alpha \in \mathcal{A}$ with $\mathcal{R}_{L_2,P}(\alpha) - \mathcal{R}_{L_2,P}^* \leq \delta$ and $A := \{x \in X : \mathcal{C}_{1,x}(\alpha) \geq \varepsilon\}$ the above considerations thus show

$$\begin{aligned} \mathcal{R}_{L_1,P}(\alpha) - \mathcal{R}_{L_1,P}^* &= \int_{\mathcal{C}_{1,x}(\alpha) \geq \varepsilon} \mathcal{C}_{1,x}(\alpha) dP_X(x) + \int_{\mathcal{C}_{1,x}(\alpha) < \varepsilon} \mathcal{C}_{1,x}(\alpha) dP_X(x) \leq \int_A b(x) dP_X(x) + \varepsilon \\ &\leq 2\varepsilon. \end{aligned}$$

From this we easily get the assertion. \blacksquare

Proof of Lemma 2.9: Let us first assume $\mathcal{C}_{L_2,Q,x}^* = \infty$. Then we have $\delta_{\max}(\varepsilon, Q, x) = \infty$ and hence *ii*) is trivially satisfied. Furthermore we have $\mathcal{M}_{L_2,Q,x}(\delta_{\max}(\varepsilon, Q, x)) = \emptyset$, and hence we obtain *i*). Let us now assume $\mathcal{C}_{L_2,Q,x}^* < \infty$. Then for $\alpha \in \mathcal{M}_{L_2,Q,x}(\delta_{\max}(\varepsilon, Q, x))$ we have

$$\mathcal{C}_{L_2,Q,x}(\alpha) - \mathcal{C}_{L_2,Q,x}^* < \delta_{\max}(\varepsilon, Q, x) = \inf_{\substack{\alpha' \in \mathcal{A} \\ \alpha' \notin \mathcal{M}_{L_1,Q,x}(\varepsilon)}} \mathcal{C}_{L_2,Q,x}(\alpha') - \mathcal{C}_{L_2,Q,x}^*,$$

which shows $\alpha \in \mathcal{M}_{L_1,Q,x}(\varepsilon)$. For the proof of the second assertion let us fix a δ with $\delta > \delta_{\max}(\varepsilon, Q, x)$. By definition this means

$$\inf_{\substack{\alpha \in \mathcal{A} \\ \alpha \notin \mathcal{M}_{L_1,Q,x}(\varepsilon)}} \mathcal{C}_{L_2,Q,x}(\alpha) - \mathcal{C}_{L_2,Q,x}^* = \delta_{\max}(\varepsilon, Q, x) < \delta,$$

and hence there exists a $\alpha \in \mathcal{M}_{L_2,Q,x}(\delta)$ with $\alpha \notin \mathcal{M}_{L_1,Q,x}(\varepsilon)$. Finally, in order to establish (7) let us fix a $\alpha \in \mathcal{A}$. Then for $\varepsilon := \mathcal{C}_{L_1,Q,x}(\alpha) - \mathcal{C}_{L_1,Q,x}^*$ we have $\alpha \notin \mathcal{M}_{L_1,Q,x}(\varepsilon)$ which implies $\alpha \notin \mathcal{M}_{L_2,Q,x}(\delta_{\max}(\varepsilon, Q, x))$ by *i*). Since the latter means

$$\mathcal{C}_{L_2,Q,x}(\alpha) \geq \mathcal{C}_{L_2,Q,x}^* + \delta_{\max}(\varepsilon, Q, x) = \mathcal{C}_{L_2,Q,x}^* + \delta_{\max}(\mathcal{C}_{L_1,Q,x}(\alpha) - \mathcal{C}_{L_1,Q,x}^*, Q, x)$$

we obtain the assertion. \blacksquare

Proof of Lemma 2.11: Since $\mathcal{M}_{L_1,Q,x}(\varepsilon)$ is an interval and $\mathcal{C}_{L_2,Q,x}(\cdot) : \mathbb{R} \rightarrow [0, \infty)$ is continuous we obtain

$$\delta_{\max}(\varepsilon, Q, x) = \min \left\{ \inf_{\alpha \leq \inf \mathcal{M}_{L_1,Q,x}(\varepsilon)} \mathcal{C}_{L_2,Q,x}(\alpha), \inf_{\alpha \geq \sup \mathcal{M}_{L_1,Q,x}(\varepsilon)} \mathcal{C}_{L_2,Q,x}(\alpha) \right\} - \mathcal{C}_{L_2,Q,x}^*.$$

Moreover, for $\alpha \in \mathbb{R}$ with $\alpha \notin \mathcal{M}_{L_1,Q,x}(\varepsilon)$ we have $\alpha \notin \mathcal{M}_{L_1,Q,x}(0^+)$, and consequently we find $\alpha \notin \mathcal{M}_{L_2,Q,x}(0^+)$. Now, it is easy to check that the map $\alpha \mapsto \mathcal{C}_{L_2,Q,x}(\alpha)$ is convex, and consequently it is strictly decreasing on $(-\infty, \inf \mathcal{M}_{L_2,Q,x}(0^+)]$ and strictly increasing on $[\sup \mathcal{M}_{L_2,Q,x}(0^+), \infty)$. Combining all observations we obtain the assertion. \blacksquare

Proof of Theorem 2.13: Let us use the shorthands $\mathcal{C}_{1,x}(\alpha)$ and $\mathcal{C}_{2,x}(\alpha)$ defined in the proof of Theorem 2.8. Assumption (10) together with Lemma 2.9 and $\mathcal{R}_{L_1,P}^* < \infty$, $\mathcal{R}_{L_2,P}^* < \infty$ then gives

$$\delta(\mathcal{C}_{1,x}(\alpha)) \leq \mathcal{C}_{2,x}(\alpha) \quad (53)$$

for all $x \in \tilde{X}$ and all $\alpha \in \mathcal{A}$. For $\alpha \in \mathcal{A}$ with $\mathcal{R}_{L_1,P}(\alpha) < \infty$ Jensen's inequality together with the definition of B_α , $\delta_{B_\alpha}^{**}(\cdot) \leq \delta(\cdot)$, and (53), now yields

$$\delta_{B_\alpha}^{**}(\mathcal{R}_{L_1,P}(\alpha) - \mathcal{R}_{L_1,P}^*) \leq \int_X \delta_{B_\alpha}^{**}(\mathcal{C}_{1,x}(\alpha)) dP_X(x) \leq \int_X \mathcal{C}_{2,x}(\alpha) dP_X(x) = \mathcal{R}_{L_2,P}(\alpha) - \mathcal{R}_{L_2,P}^*.$$

Finally for $\alpha \in \mathcal{A}$ with $\mathcal{R}_{L_1,P}(\alpha) = \infty$ we have $B_\alpha = \infty$. If $\delta_\infty^{**}(\infty) = 0$ there is nothing to prove, and hence let us assume $\delta_\infty^{**}(\infty) > 0$. Then observe that because of $0 = \delta(0) = \delta_\infty^{**}(0)$ and its convexity the function δ^{**} is increasing. Consequently, if there is a $t_0 > 0$ with $\delta_\infty^{**}(t_0) = \infty$, there exist constants $c_1, c_2 \in (0, \infty)$ with $t \leq c_1 \delta_\infty^{**}(t) + c_2$ for all $t \in [0, \infty]$. Furthermore, if δ^{**} is finite on $[0, \infty)$ then there exists a $t_0 \geq 0$ and a $c_0 > 0$ such that the (Lebesgue)-almost surely defined derivative of δ_∞^{**} satisfies $(\delta_\infty^{**})'(t) \geq c_0$ for almost all $t \geq t_0$. By Lebesgue's version of the fundamental theorem of calculus (see e.g. the Theorems 26-28 in Chapter X of [10] or the Theorems 271, 269, and 274 in [15]) we then also find constants $c_1, c_2 \in (0, \infty)$ with $t \leq c_1 \delta_\infty^{**}(t) + c_2$ for all $t \in [0, \infty]$. This together with (53) yields

$$\infty = \int_X \mathcal{C}_{1,x}(\alpha) dP_X(x) \leq c_1 \int_X \delta_\infty^{**}(\mathcal{C}_{1,x}(\alpha)) dP_X(x) + c_2 \leq c_1 (\mathcal{R}_{L_2,P}(\alpha) - \mathcal{R}_{L_2,P}^*) + c_2$$

and hence we have $\mathcal{R}_{L_2,P}(\alpha) - \mathcal{R}_{L_2,P}^* = \infty$. This shows the assertion. \blacksquare

Proof of Lemma 2.16: Let us fix an $\varepsilon > 0$. If $\delta_{\max}(\varepsilon, \mathcal{Q}) = 0$ there is nothing to prove and hence we assume $\delta_{\max}(\varepsilon, \mathcal{Q}) > 0$ without loss of generality. Then there exists a strictly positive sequence (δ_n) with $\delta_n \nearrow \delta_{\max}(\varepsilon, \mathcal{Q})$ for $n \rightarrow \infty$, and $\mathcal{M}_{L_2,Q,x}(\delta_n) \subset \mathcal{M}_{L_1,Q,x}(\varepsilon)$ for all $n \in \mathbb{N}$. Now let us fix an $\alpha \in \mathcal{M}_{L_2,Q,x}(\delta_{\max}(\varepsilon, \mathcal{Q}))$. Then we have $\mathcal{C}_{L_2,Q,x}(\alpha) - \mathcal{C}_{L_2,Q,x}^* < \delta_{\max}(\varepsilon, \mathcal{Q})$ and hence there exists an $n \in \mathbb{N}$ with $\mathcal{C}_{L_2,Q,x}(\alpha) - \mathcal{C}_{L_2,Q,x}^* < \delta_n$. Obviously this implies $\alpha \in \mathcal{M}_{L_2,Q,x}(\delta_n) \subset \mathcal{M}_{L_1,Q,x}(\varepsilon)$, i.e. we have shown the assertion. \blacksquare

Proof of Lemma 2.17: Let us write $\delta(\varepsilon) := \inf_{x \in X} \delta_{\max}(\varepsilon, Q, x)$ for $\varepsilon > 0$. Since $\delta(\varepsilon) \leq \delta_{\max}(\varepsilon, Q, x)$ for all $x \in X$, $Q \in \mathcal{Q}$, and $\varepsilon > 0$, we then obtain

$$\mathcal{M}_{L_2,Q,x}(\delta(\varepsilon)) \subset \mathcal{M}_{L_2,Q,x}(\delta_{\max}(\varepsilon, Q, x)) \subset \mathcal{M}_{L_1,Q,x}(\varepsilon).$$

This shows $\delta(\varepsilon) \leq \delta_{\max}(\varepsilon, \mathcal{Q})$. In order to prove the converse inequality let us assume $\delta(\varepsilon) < \delta_{\max}(\varepsilon, \mathcal{Q})$ for some $\varepsilon > 0$. Then there exist $Q \in \mathcal{Q}$ and $x \in X$ with $\delta_{\max}(\varepsilon, Q, x) < \delta_{\max}(\varepsilon)$. However, Lemma 2.16 shows $\mathcal{M}_{L_2,Q,x}(\delta_{\max}(\varepsilon, \mathcal{Q})) \subset \mathcal{M}_{L_1,Q,x}(\varepsilon)$, and hence we find a contradiction by *ii*) in Lemma 2.9. \blacksquare

Proof of Theorem 2.18: Let us fix an $x \in X$ and a $Q \in \mathcal{Q}$. Furthermore, let P be the distribution on $X \times Y$ with $P_X = \delta_{\{x\}}$, where $\delta_{\{x\}}$ is the Dirac measure in x , and $P(\cdot|x) = Q$. Then P is of type \mathcal{Q} and we have $\mathcal{R}_{L_i,P}(\alpha) = \mathcal{C}_{L_i,Q,x}(\alpha)$ and $\mathcal{R}_{L_i,P}^* = \mathcal{C}_{L_i,Q,x}^* < \infty$ for $i = 1, 2$ and all measurable $\alpha \in \mathcal{A}$. Consequently our assumption yields

$$\delta(\mathcal{C}_{L_1,Q,x}(\alpha) - \mathcal{C}_{L_1,Q,x}^*) \leq \mathcal{C}_{L_2,Q,x}(\alpha) - \mathcal{C}_{L_2,Q,x}^*$$

for all $\alpha \in \mathcal{A}$. Now let $\varepsilon > 0$ and $\alpha \in \mathcal{M}_{L_2,Q,x}(\delta(\varepsilon))$. Then we have $\mathcal{C}_{L_2,Q,x}(\alpha) - \mathcal{C}_{L_2,Q,x}^* < \delta(\varepsilon)$ and hence the above inequality yields $\delta(\mathcal{C}_{L_1,Q,x}(\alpha) - \mathcal{C}_{L_1,Q,x}^*) < \delta(\varepsilon)$. Since δ is monotonically increasing the latter shows $\mathcal{C}_{L_1,Q,x}(\alpha) - \mathcal{C}_{L_1,Q,x}^* < \varepsilon$, i.e. we have found $\mathcal{M}_{L_2,Q,x}(\delta(\varepsilon)) \subset \mathcal{M}_{L_1,Q,x}(\varepsilon)$. Lemma 2.9 then shows $\delta(\varepsilon) \leq \delta_{\max}(\varepsilon, Q, x)$, and hence L_2 is uniformly L_1 -calibrated with respect to \mathcal{Q} . \blacksquare

Proof of Theorem 2.19: Let us use the shorthands $\mathcal{C}_{1,x}(\alpha)$ and $\mathcal{C}_{2,x}(\alpha)$ defined in the proof of Theorem 2.8. Then our assumption on $\delta_{\max}(\varepsilon, x)$ together with (7) yields $(\mathcal{C}_{2,x}(\alpha))^q b(x) \leq \mathcal{C}_{1,x}(\alpha)$ for all $x \in X$, $\alpha \in \mathcal{A}$, and hence using Hölder's inequality with r defined by $q = \frac{1}{r} + 1$ gives

$$\mathcal{R}_{L_1, P}(\alpha) - \mathcal{R}_{L_1, P}^* \leq \left(\int_X (b(x))^{-\frac{1}{q}} (\mathcal{C}_{1,x}(\alpha))^{\frac{1}{q}} dP_X(x) \right)^{\frac{q}{q-1}} \leq \left(\int_X b^{-r} dP_X \right)^{\frac{1}{q-1}} \left(\int_X \mathcal{C}_{1,x}(\alpha) dP_X(x) \right)^{\frac{1}{q-1}}.$$

Since $p \geq r$ we then find the assertion. ■

For the proof of Theorem 2.22 we need a method for finding measurable selections from multi-valued maps of specific forms. We begin with some basics: Let X , Y , and Z be measurable spaces, $h : X \times Z \rightarrow Y$ and $A \subset Y$ be measurable, and

$$\begin{aligned} F : X &\rightarrow 2^Z \\ x &\mapsto \{z \in Z : h(x, z) \in A\}. \end{aligned} \quad (54)$$

Furthermore we write

$$\begin{aligned} \text{Dom } F &:= \{x \in X : F(x) \neq \emptyset\}, \\ \text{Gr } F &:= \{(x, z) \in X \times Z : z \in F(x)\}, \\ F^{-1}(B) &:= \{x \in X : F(x) \cap B \neq \emptyset\} \quad B \subset Z. \end{aligned}$$

Note that we have $\text{Dom } F = F^{-1}(Z)$ and $\text{Gr } F = \{(x, z) \in X \times Z : h(x, z) \in A\}$, so that in particular $\text{Gr } F$ is measurable. Furthermore, if $\pi_X : X \times Z \rightarrow X$ denotes the projection onto X then we have

$$\text{Dom } F = \{x \in X : \exists z \in Z \text{ with } h(x, z) \in A\} = \pi_X \left(\{(x, z) \in X \times Z : h(x, z) \in A\} \right) = \pi_X(\text{Gr } F).$$

Now the following result provides a sufficient condition under which $\text{Dom } F$ is measurable and F admits measurable selections.

Lemma 4.1 *Let (X, \mathcal{X}) be a complete measurable space, Z be a Polish space equipped with its Borel σ -algebra, and Y be a measurable space. Furthermore let $h : X \times Z \rightarrow Y$ be a measurable map, $A \subset Y$ be measurable and $F : X \rightarrow 2^Z$ be defined by (54). Then the following are true:*

- i) $\text{Dom } F$ is measurable.
- ii) There exists a sequence of measurable functions $f_n : X \rightarrow Z$ such that for all $x \in \text{Dom } F$ the set $\{f_n(x) : n \in \mathbb{N}\}$ is dense in $F(x)$.
- iii) Let $\varphi : X \times Z \rightarrow [0, \infty]$ be measurable and $\psi : X \rightarrow [0, \infty]$ be defined by

$$\psi(x) := \inf_{z \in F(x)} \varphi(x, z), \quad x \in X. \quad (55)$$

Then ψ is measurable. Furthermore, for all $n \geq 1$ there exists a measurable $f_n : X \rightarrow Z$ such that for all $x \in \text{Dom } F$ we have $f_n(x) \in F(x)$ and $\varphi(x, f_n(x)) \leq \psi(x) + 1/n$, and consequently

$$\psi(x) = \inf_{n \in \mathbb{N}} \varphi(x, f_n(x)).$$

In addition, if the infimum in (55) is attained for all $x \in \text{Dom } X$ then there exists a measurable function $f : X \rightarrow Z$ with $\psi(x) = \varphi(x, f(x))$ for all $x \in \text{Dom } F$.

Proof: *i*). Let us first recall that the so-called projection theorem [7, Thm. III.23, p. 75] ensures $\pi_X(B) \in \mathcal{X}$ for all $B \in \mathcal{X} \otimes \mathcal{B}(Z)$. Now the assertion directly follows from $\text{Dom } F = \pi_X(\text{Gr } F)$ and the projection theorem.

ii). Let $\tilde{\mathcal{X}}$ be the trace σ -algebra of \mathcal{X} on $\text{Dom } F$. Then it is easy to see that $\tilde{\mathcal{X}}$ is a complete σ -algebra. Furthermore, $F|_{\text{Dom } F} : \text{Dom } F \rightarrow 2^Z$ obviously maps $\text{Dom } F$ to non-empty subsets of Z and in addition we have

$$\text{Gr}(F|_{\text{Dom } F}) = \{(x, z) \in \text{Dom } F \times Z : z \in F(x)\} = \text{Gr } F \cap (\text{Dom } F \times Z),$$

and hence $\text{Gr}(F|_{\text{Dom } F})$ is measurable. Now Aumann's selection theorem in the form of [7, Thm. III.22, p. 74] gives a sequence of $\tilde{\mathcal{X}}$ -measurable functions $\tilde{f}_n : \text{Dom } F \rightarrow \mathcal{T}$ such that $\{f_n(x) : n \in \mathbb{N}\}$ is dense in $F(x)$ for all $x \in \text{Dom } F$. Extending these functions to measurable functions $f_n : X \rightarrow \mathcal{T}$ gives the assertion.

iii) The measurability follows from [7, Lem. III.39, p. 86]. Furthermore, on the measurable set $\{x \in X : \psi(x) = \infty\}$ there is nothing to prove and hence we may restrict our considerations to $\text{Dom } F \cap \{x \in X : \psi(x) < \infty\}$ equipped with the trace σ -algebra of \mathcal{X} . Then the existence of f_n is shown in [7, p. 87]. Finally, the existence of a measurable $f : X \rightarrow Z$ is shown in [7, p. 86]. ■

Proof of Theorem 2.22: First note that the measurability of $(x, t) \mapsto \mathcal{C}_{L, P(\cdot|x), x}(t)$ can be shown using standard arguments.

i). If $\mathcal{C}_{L, P(\cdot|x), x}^* = \infty$ for some $x \in X$, then (2) cannot be true for this x , and hence \hat{L} is not P -minimizable. To see the converse implication observe that the multivalued map $F : X \rightarrow 2^{\mathcal{T}}$, defined by $F(x) := \mathcal{T}$ for all $x \in X$, is obviously of the form (54) for arbitrary measurable h and $A := \mathcal{T}$. We write $\varphi(x, t) := \mathcal{C}_{L, P(\cdot|x), x}(t)$ for all $x \in X, t \in \mathcal{T}$, so that $\varphi : X \times \mathcal{T} \rightarrow [0, \infty]$ is measurable. Then we have

$$\mathcal{C}_{L, P(\cdot|x), x}^* = \inf_{t \in F(x)} \varphi(x, t) \quad (56)$$

for all $x \in X$, and consequently Lemma 4.1 shows that for all $n \geq 1$ there exists a measurable function $f_n : X \rightarrow \mathcal{T}$ with

$$\mathcal{C}_{L, P(\cdot|x), x}(f_n(x)) = \varphi(x, f_n(x)) \leq \mathcal{C}_{L, P(\cdot|x), x}^* + 1/n < \mathcal{C}_{L, P(\cdot|x), x}^* + 2/n$$

for all $x \in \text{Dom } F = X$. From this we easily conclude that \hat{L} is P -minimizable.

ii). This assertion also follows from Lemma 4.1 since $\mathcal{M}_{L, P(\cdot|x), x}(0^+) \neq \emptyset$ for all $x \in \mathcal{X}$ ensures that the infimum in (56) is attained for all $x \in X$. ■

Proof of Theorem 2.23: Without loss of generality we may assume $\mathcal{C}_{L_i, P(\cdot|x), x}^* < \infty$ for all $x \in X$ and $i = 1, 2$. To show the measurability of $x \mapsto \delta_{\max}(\varepsilon, P(\cdot|x), x)$ we equip $[0, \infty]$ with the Borel σ -algebra, write $A := [\varepsilon, \infty]$, and define $h : X \times \mathcal{T} \rightarrow [0, \infty]$ by

$$h(x, t) := \mathcal{C}_{L_1, P(\cdot|x), x}(t) - \mathcal{C}_{L_1, P(\cdot|x), x}^*, \quad (x, t) \in X \times \mathcal{T}.$$

Then h is measurable and for $F : X \rightarrow 2^{\mathcal{T}}$ defined by (54) we have $\mathcal{T} \setminus \mathcal{M}_{L_1, P(\cdot|x), x}(\varepsilon) = F(x)$ for all $x \in X$. Furthermore, $\varphi : X \times \mathcal{T} \rightarrow [0, \infty]$ defined by

$$\varphi(x, t) := \mathcal{C}_{L_2, P(\cdot|x), x}(t) - \mathcal{C}_{L_2, P(\cdot|x), x}^*, \quad (x, t) \in X \times \mathcal{T},$$

is also measurable. For all $x \in X$ our construction yields

$$\delta_{\max}(\varepsilon, P(\cdot|x), x) = \inf_{t \in F(x)} \varphi(x, t),$$

and consequently we obtain the measurability of $x \mapsto \delta_{\max}(\varepsilon, P(\cdot|x), x)$ by Lemma 4.1.

Let us now assume that there exists an $\varepsilon > 0$ such that $B := \{x \in X : \delta_{\max}(\varepsilon, x) = 0\}$ satisfies $P_X(B) > 0$. With the above notations we obviously have $B \subset \text{Dom } F$. Moreover for all $n \geq 1$ Lemma 4.1 gives us a measurable function $f_n^{(1)} : X \rightarrow \mathcal{T}$ with

$$\mathcal{C}_{L_2, P(\cdot|x), x}(f_n^{(1)}(x)) - \mathcal{C}_{L_2, P(\cdot|x), x}^* \leq \frac{1}{n} \quad \text{and} \quad \mathcal{C}_{L_1, P(\cdot|x), x}(f_n^{(1)}) - \mathcal{C}_{L_1, P(\cdot|x), x}^* \geq \varepsilon$$

for all $x \in B$. Furthermore, since L_2 is P -minimizable there also exist measurable functions $f_n^{(2)} : X \rightarrow \mathcal{T}$ with

$$\mathcal{C}_{L_2, P(\cdot|x), x}(f_n^{(2)}(x)) - \mathcal{C}_{L_2, P(\cdot|x), x}^* \leq \frac{1}{n}$$

for all $x \in X$. We define $f_n : X \rightarrow \mathcal{T}$ by $f_n(x) := f_n^{(1)}(x)$ if $x \in B$, and $f_n(x) := f_n^{(2)}(x)$ otherwise. Then f_n is measurable and our construction yields both $\lim_{n \rightarrow \infty} \mathcal{R}_{L_2, P}(f_n) = \mathcal{R}_{L_2, P}^*$ and

$$\mathcal{R}_{L_1, P}(f) - \mathcal{R}_{L_1, P}^* \geq \int_B \mathcal{C}_{L_1, P(\cdot|x), x}(t) - \mathcal{C}_{L_1, P(\cdot|x), x}^* dP_X(x) \geq \varepsilon P_X(B).$$

From this we easily obtain the assertion. ■

Proof of Theorem 2.26: Since $\delta_{\max}(\cdot, \mathcal{Q})$ is monotonously increasing, the set

$$U := \{\varepsilon > 0 : \delta_{\max}(\cdot, \mathcal{Q}) \text{ not continuous in } \varepsilon\}$$

is at most countable, and hence there exists a sequence (ε_n) with $\{\varepsilon_n : n \in \mathbb{N}\} = U \cup \{r \in \mathcal{Q} : r > 0\}$. Then for all $n, m \geq 1$ there exists a distribution $Q_{n,m} \in \mathcal{Q}$ with

$$\frac{1}{m} + \delta_{\max}(\varepsilon_n, \mathcal{Q}) > \delta_{\max}(\varepsilon_n, Q_{n,m}) = \inf_{t \in \mathcal{T}} \mathcal{C}_{L_2, Q_{n,m}}(t) - \mathcal{C}_{L_2, Q_{n,m}}^*.$$

$t \notin \mathcal{M}_{L_1, Q_{n,m}}(\varepsilon_n)$

Therefore for all $n, m \geq 1$ there exist $t_{n,m} \in \mathcal{T}$ with $t_{n,m} \notin \mathcal{M}_{L_1, Q_{n,m}}(\varepsilon_n)$ and

$$\mathcal{C}_{L_2, Q_{n,m}}(t_{n,m}) - \mathcal{C}_{L_2, Q_{n,m}}^* < \delta_{\max}(\varepsilon_n, \mathcal{Q}) + \frac{1}{m}.$$

Now, let $A_{n,m} \subset X$, $n \geq 1$, be according to our assumption. Note that without loss of generality we can additionally assume $X = \bigcup_{n,m} A_{n,m}$. Consequently, for $x \in X$ there exists a unique $(n_x, m_x) \in \mathbb{N} \times \mathbb{N}$ with $x \in A_{n_x, m_x}$. Furthermore, let P be the distribution on $X \times Y$ which is defined by the conditions $P_X = \mu$ and $P(\cdot|x) = Q_{n_x, m_x}$, $x \in X$. Now let us fix an $n \geq 1$. Then for all $x \in \bigcup_m A_{n,m}$ our construction gives $t_{n, m_x} \notin \mathcal{M}_{L_1, P(\cdot|x), x}(\varepsilon_n)$, and hence we obtain

$$\delta_{\max}(\varepsilon_n, P) \leq \mathcal{C}_{L_2, Q_{n, m_x}}(t_{n, m_x}) - \mathcal{C}_{L_2, Q_{n, m_x}}^* < \delta_{\max}(\varepsilon_n, \mathcal{Q}) + \frac{1}{m_x}.$$

Minimizing over $x \in \bigcup_m A_{n,m}$ then gives $\delta_{\max}(\varepsilon_n, P) \leq \delta_{\max}(\varepsilon_n, \mathcal{Q})$ for all $n \in \mathbb{N}$, and consequently we have $\delta_{\max}(\varepsilon_n, P) = \delta_{\max}(\varepsilon_n, \mathcal{Q})$ for all $n \in \mathbb{N}$. In particular, this shows $\delta_{\max}(\varepsilon, P) = \delta_{\max}(\varepsilon, \mathcal{Q})$ for all $\varepsilon \in U$. Now let $\varepsilon > 0$ with $\varepsilon \notin U$. There there exists a sub-sequence (ε_{n_k}) of (ε_n) with $\varepsilon_{n_k} \searrow \varepsilon$ for $k \rightarrow \infty$. This gives $\delta_{\max}(\varepsilon_{n_k}, P) = \delta_{\max}(\varepsilon_{n_k}, \mathcal{Q}) \rightarrow \delta_{\max}(\varepsilon, \mathcal{Q})$, and hence we have

$$\delta_{\max}(\varepsilon, \mathcal{Q}) \geq \inf_{k \geq 1} \delta_{\max}(\varepsilon_{n_k}, P) \geq \inf_{\varepsilon' \geq \varepsilon} \delta_{\max}(\varepsilon', P) = \delta_{\max}(\varepsilon, P)$$

by the monotonicity of $\delta_{\max}(\cdot, P)$. From this we easily obtain (15). ■

Proof of Theorem 2.30: To shorten notations we write

$$\begin{aligned}\mathcal{C}_{1,x}(f) &:= \mathcal{C}_{L_1,P(\cdot|x),x}(f(x)) - \mathcal{C}_{L_1,P(\cdot|x),x}^*, & \text{and} \\ \mathcal{C}_{2,x}(f) &:= \mathcal{C}_{L_2,P(\cdot|x),x}(f(x)) - \mathcal{C}_{L_2,P(\cdot|x),x}^*\end{aligned}$$

for $x \in X$ and measurable $f : X \rightarrow \mathcal{T}$. Furthermore, for $s > 0$ we write

$$C(s) := \{x \in X : A(x) \neq \mathcal{T}, \text{ and } \delta_{\max}(h(x), P(\cdot|x), x) \geq s h(x)\}.$$

By (7) and (17) we then obtain

$$\begin{aligned}\mathcal{R}_{L_1,P}(f) - \mathcal{R}_{L_1,P}^* &= \int_{B(s)} \mathbf{1}_A(x, f(x)) h(x) dP_X(x) + \int_{C(s)} \mathbf{1}_A(x, f(x)) h(x) dP_X(x) \\ &\leq \int \mathbf{1}_{B(s)} h dP_X + s^{-1} \int_{C(s)} \delta_{\max}(h(x), P(\cdot|x), x) \mathbf{1}_A(x, f(x)) dP_X(x) \\ &\leq (cs)^\alpha + s^{-1} \int_{C(s)} \delta_{\max}(\mathcal{C}_{1,x}(f), P(\cdot|x), x) dP_X(x) \\ &\leq (cs)^\alpha + s^{-1} (\mathcal{R}_{L_2,P}(f) - \mathcal{R}_{L_2,P}^*).\end{aligned}$$

For $\alpha < \infty$, the assertion then follows by setting $s := (\alpha c^\alpha)^{-\frac{1}{\alpha+1}} (\mathcal{R}_{L_2,P}(f) - \mathcal{R}_{L_2,P}^*)^{\frac{1}{\alpha+1}}$ and using $\alpha^{-\frac{\alpha}{\alpha+1}} + \alpha^{\frac{1}{\alpha+1}} \leq 2$.

Furthermore, for $\alpha = \infty$ the assertion follows by setting $s^{-1} := 2c$. ■

Proof of Lemma 2.34: Let P be a distribution on $X \times Y$ with $P(\cdot|x) \in \mathcal{Q}_{\min}(L)$ for all $x \in X$. We write $\bar{X} := X \times \mathbb{R}$ and $Z := \mathbb{R}$. Furthermore, for $\bar{x} = (x, t) \in \bar{X}$ and $t' \in Z$ we define

$$\begin{aligned}h(\bar{x}, t') &:= \mathcal{C}_{L,P(\cdot|x)}(t') - \mathcal{C}_{L,P(\cdot|x)}^*, \\ F(\bar{x}) &:= \{t' \in \mathbb{R} : h(\bar{x}, t') = 0\}, & \text{and} \\ \varphi(\bar{x}, t') &:= |t - t'|.\end{aligned}$$

For the P -instance \check{L}_P of \check{L} we then have

$$\check{L}_P(x, t) = \inf_{t' \in \mathcal{M}_{L,P(\cdot|x)}(0^+)} |t - t'| = \inf_{t' \in F(\bar{x})} \varphi(\bar{x}, t'),$$

and consequently we obtain the assertion by part *iii*) of Lemma 4.1. ■

Proof of Theorem 2.37: For a fixed $\rho > 0$ we write $A_\rho = \{(Q, t) \in \mathcal{Q} \times \mathbb{R} : \check{L}(Q, t) \geq \rho\}$. By Lemma 2.34 we then see that $\bar{L} := \mathbf{1}_{A_\rho}$ defines a template loss function, whose P -instance \bar{L}_P is a detection loss with $h = \mathbf{1}_X$. Furthermore, we have $\mathcal{M}_{\bar{L},Q}(\varepsilon) = \{t \in \mathbb{R} : \check{L}(Q, t) \geq \rho\} = \mathcal{M}_{\check{L},Q}(\rho)$ for all $\varepsilon > 0$ and $Q \in \mathcal{Q}$, and thus we find

$$\delta_{\max, \bar{L}, L}(\varepsilon, Q) = \delta_{\max, \check{L}, L}(\rho, Q) > 0.$$

In other words, L is \bar{L} -calibrated with respect to \mathcal{Q} . For $\varepsilon > 0$ Theorem 2.29 thus gives a $\delta > 0$ such that for $f : X \rightarrow \mathbb{R}$ with $\mathcal{R}_{L,P}(f) < \mathcal{R}_{L,P}^* + \delta$ we have

$$P_X(\{x \in X : \check{L}_P(x, f(x)) \geq \rho\}) = \mathcal{R}_{\bar{L}_P,P}(f) - \mathcal{R}_{\bar{L}_P,P}^* < \varepsilon.$$

■

Proof of Proposition 2.40: We write $\alpha := \frac{pq}{p+1}$ and $\bar{L} := \check{L}^\alpha$. Then \bar{L} is a template loss function by Lemma 2.34, and since $\mathcal{M}_{\bar{L},Q}(\varepsilon) = \{t \in \mathbb{R} : \bar{L}(Q, t) < \varepsilon\}$ we easily find

$$\delta_{\max, \bar{L}, L}(\varepsilon, Q) = \delta_{\max, \check{L}, L}(\varepsilon^{1/\alpha}, Q), \quad \varepsilon > 0, Q \in \mathcal{Q}.$$

Furthermore, we obviously have $\frac{q}{\alpha} \geq \frac{p+1}{p}$ and $\mathcal{R}_{\bar{L},P}^* = 0$, and therefore Theorem 2.19 yields

$$\|x \mapsto \check{L}_P(x, f(x))\|_\alpha^\alpha = \mathcal{R}_{\bar{L},P}(f) - \mathcal{R}_{\bar{L},P}^* \leq \|b^{-1}\|_p^{\frac{\alpha}{q}} \left(\mathcal{R}_{L,P}(f) - \mathcal{R}_{L,P}^* \right)^{\frac{\alpha}{q}}.$$

■

Appendix

In this appendix we discuss some simple properties of convex functions. Recall that a function $f : I \rightarrow \mathbb{R}$ on an interval $I \subset \mathbb{R}$ is called *convex* if for all $x_1, x_2 \in I$ and all $\alpha \in [0, 1]$ we have

$$f(\alpha x_1 + (1 - \alpha)x_2) \leq \alpha f(x_1) + (1 - \alpha)f(x_2).$$

Furthermore, it is called *strictly convex* if this inequality is strict for all $x_1, x_2 \in I$ with $x_1 \neq x_2$, and all $\alpha \in (0, 1)$. Obviously, every strictly convex function is also convex. Furthermore, it is well-known that convex functions $f : I \rightarrow \mathbb{R}$ are continuous on the interior of I . The following lemma describes some less trivial relations between the different notions of convexity:

Lemma A.1 *Let $f : I \rightarrow \mathbb{R}$ be a function. Then we have:*

- i) If f is convex and satisfies $f(\alpha_0 x_1 + (1 - \alpha_0)x_2) = \alpha_0 f(x_1) + (1 - \alpha_0)f(x_2)$ for some $x_1, x_2 \in I$, $\alpha_0 \in [0, 1]$, then actually for all $\alpha \in [0, 1]$ we have*

$$f(\alpha x_1 + (1 - \alpha)x_2) = \alpha f(x_1) + (1 - \alpha)f(x_2). \quad (57)$$

- ii) If f is continuous then f is convex if and only if for all $x_1, x_2 \in I$ we have*

$$f\left(\frac{x_1 + x_2}{2}\right) \leq \frac{1}{2}f(x_1) + \frac{1}{2}f(x_2). \quad (58)$$

- iii) If f is continuous then f is strictly convex if and only if for all $x_1, x_2 \in I$ with $x_1 \neq x_2$ we have*

$$f\left(\frac{x_1 + x_2}{2}\right) < \frac{1}{2}f(x_1) + \frac{1}{2}f(x_2). \quad (59)$$

- iv) If f is uniformly convex and continuous then it is strictly convex. Conversely, if I is compact and f is strictly convex and continuous then it is actually uniformly convex.*

Proof: *i).* This assertion can be shown using elementary calculations.

ii). This follows from [2, Thm. 8 and 10].

iii). If (59) holds then we have already seen that f is convex. Consequently, if f was not strictly convex we would have (57). However, by *i)* we could then assume $\alpha_0 = \frac{1}{2}$ which would give a contradiction.

iv) The first assertion follows from *iii)* and the second assertion is trivial. ■

Our next aim is to investigate the modulus of convexity. Although this concept, in an equivalent formulation, has already been introduced in 1966 (see [22], [16]) almost nothing that is useful for us, seems to be known (see however [5, 34] and the references therein for some general information on the modulus). Therefore, we present the following two lemmas which provide some ways to simplify the computation of $\delta_f(\varepsilon)$:

Lemma A.2 *Let $\emptyset \neq I \subset \mathbb{R}$ be an interval and $f : I \rightarrow \mathbb{R}$ be strictly convex. Then we have*

$$\delta_f(2\varepsilon) = \inf \left\{ \frac{f(x-\varepsilon) + f(x+\varepsilon)}{2} - f(x) : x \text{ satisfies } x-\varepsilon \in I \text{ and } x+\varepsilon \in I \right\}, \quad \varepsilon > 0.$$

Proof: For fixed $x_1 \in I$ we define $h_{x_1} : I \rightarrow [0, \infty)$ by $h_{x_1}(x_2) := \frac{f(x_1)+f(x_2)}{2} - f(\frac{x_1+x_2}{2})$, $x_2 \in I$. The fundamental theorem of calculus for Lebesgue integrals then shows that the derivative $h'_{x_1}(x_2)$ exists for almost all x_2 , and an easy calculation shows $h'_{x_1}(x_2) = \frac{f'(x_2)}{2} - \frac{1}{2}f'(\frac{x_1+x_2}{2})$ for such x_2 . Furthermore, since f is strictly convex and h_{x_1} has a unique minimum at x_1 , we obtain $h'_{x_1}(x_2) < 0$ if $x_2 < x_1$, and $h'_{x_1}(x_2) > 0$ if $x_2 > x_1$. The fundamental theorem of calculus for Lebesgue integrals then shows that h_{x_1} is strictly decreasing on $(-\infty, x_1) \cap I$, and strictly increasing on $(x_1, \infty) \cap I$, and thus we have

$$\delta_f(2\varepsilon) = \min \left\{ \inf_{\substack{x_1 \in I \\ x_1-2\varepsilon \in I}} h_{x_1}(x_1-2\varepsilon), \inf_{\substack{x_1 \in I \\ x_1+2\varepsilon \in I}} h_{x_1}(x_1+2\varepsilon) \right\} = \inf_{\substack{x_1+\varepsilon \in I \\ x_1-\varepsilon \in I}} h_{x_1-\varepsilon}(x_1+\varepsilon),$$

where in the last step we used $h_{x_1-\varepsilon}(x_1+\varepsilon) = h_{x_1+\varepsilon}(x_1-\varepsilon)$. ■

With the help of the following lemma we can often estimate the modulus of convexity.

Lemma A.3 *Let $I \subset \mathbb{R}$ be a symmetric interval, i.e. $x \in I$ implies $-x \in I$. Then for all strictly convex, symmetric $f : I \rightarrow [0, \infty)$ and all $\varepsilon > 0$ we have*

$$\delta_f(2\varepsilon) = \inf_{\substack{x \geq 0 \\ x+\varepsilon \in I}} \frac{f(x-\varepsilon) + f(x+\varepsilon)}{2} - f(x) = \frac{1}{2} \inf_{\substack{x \geq 0 \\ x+\varepsilon \in I}} \int_x^{x+\varepsilon} (f'(t) - f'(t-\varepsilon)) dt.$$

Furthermore, if $I = \mathbb{R}$, then for all $x \geq 12\varepsilon$ we have

$$f(x) \geq \frac{2\delta_f(2\varepsilon)x^2}{\varepsilon^2}.$$

Proof: The first equation follows from Lemma A.2 and the symmetry assumptions. Furthermore, by the fundamental theorem of calculus for Lebesgue integrals we obtain

$$f(x+\varepsilon) + f(x-\varepsilon) = 2f(x) + \int_x^{x+\varepsilon} (f'(t) - f'(t-\varepsilon)) dt, \quad (60)$$

and hence the second equation follows. Finally, in order to show the last assertion we first observe that f has a minimum at 0, and hence we have $f'(t) \geq 0$ for all $t \geq 0$ where the derivative exists. We write $b := 2\delta_f(2\varepsilon)$, and $x_n := 2\varepsilon n$ for $n \geq 1$. These definitions together with (60) yield real numbers $t_n \in [x_n, x_n + \varepsilon]$, $n \geq 1$, that satisfy

$$b \leq \int_{x_n}^{x_n+\varepsilon} (f'(t) - f'(t-\varepsilon)) dt \leq \varepsilon(f'(t_n) - f'(t_n - \varepsilon)),$$

i.e. we obtain $f'(t_n) \geq f'(t_n - \varepsilon) + \frac{b}{\varepsilon}$ for all $n \geq 1$. Furthermore we have $t_n - \varepsilon \geq x_n - \varepsilon = x_{n-1} + \varepsilon \geq t_{n-1}$ and hence $f'(t_n - \varepsilon) \geq f'(t_{n-1})$, $n \geq 2$. By induction we thus find $f'(t_{n+1}) \geq f'(t_1) + \frac{bn}{\varepsilon} \geq \frac{bn}{\varepsilon}$ for all $n \geq 1$. Now let $t \geq 6\varepsilon$ such that $f'(t)$ exists. Then there is an $n \geq 3$ with $2\varepsilon n \leq t < 2\varepsilon(n+1)$, and hence we get

$$f'(t) \geq f'(t_{n-1}) \geq \frac{b(n-2)}{\varepsilon} > \frac{b(x-6\varepsilon)}{2\varepsilon^2}.$$

Consequently, for $x \geq 12\varepsilon$ the fundamental theorem of calculus for Lebesgue integrals gives

$$f(x) = f(6\varepsilon) + \int_{6\varepsilon}^x f'(t)dt \geq \frac{b}{2\varepsilon^2} \int_{6\varepsilon}^x (t-6\varepsilon)dt = \frac{b(x-6\varepsilon)^2}{4\varepsilon^2} \geq \frac{bx^2}{\varepsilon^2}.$$

■

Example A.4 We write $I := [-B, B]$, and for $0 < p < 2$ we define $\psi : I \rightarrow [0, \infty]$ by $\psi(t) := |t|^p$, $t \in I$. Then for all $\varepsilon \in [0, B]$ we have

$$\frac{p(p-1)}{2} B^{p-2} \varepsilon^2 \leq \delta_\psi(2\varepsilon) \leq \frac{p}{2(p-1)^2} B^{p-2} \varepsilon^2.$$

To see this let $t \in [0, 1]$ and $a \in (0, 1)$. Then we have $(1-0)^a = 1-a0$ and $-a(1-t)^{a-1} \leq -a$, and thus we obtain $(1-t)^a \leq 1-at$ by the fundamental theorem of calculus. For $s := 1/t$ we hence find $(1-\frac{1}{s})^a \leq 1-\frac{a}{s}$ which is equivalent to $(s-1)^a \leq s^a - as^{a-1}$. Since in addition $s^{a-1} \leq (s-1)^{a-1}$ implies $s^a - s^{a-1} \leq (s-1)^a$ we have

$$as^{a-1} \leq s^a - (s-1)^a \leq s^{a-1} \quad (61)$$

for all $0 < a < 1$ and all $s \geq 1$. Now an easy calculation shows

$$\psi'(t) - \psi'(t-\varepsilon) = p \begin{cases} t^{p-1} - (t-\varepsilon)^{p-1} & \text{if } t \geq \varepsilon \\ t^{p-1} + (\varepsilon-t)^{p-1} & \text{if } 0 \leq t \leq \varepsilon. \end{cases}$$

Furthermore, for $s := \frac{t}{\varepsilon} \geq 1$ we have $\psi'(t) - \psi'(t-\varepsilon) = \varepsilon^{p-1}(s^{p-1} - (s-1)^{p-1})$, and thus our preliminary considerations give

$$(p-1)\varepsilon t^{p-2} = (p-1)\varepsilon^{p-1} s^{p-2} \leq \psi'(t) - \psi'(t-\varepsilon) \leq \varepsilon^{p-1} s^{p-2} = \varepsilon t^{p-2}.$$

From this we can easily prove the assertion using basic calculations.

Example A.5 We write $I := [-B, B]$ and define $\psi : I \rightarrow [0, \infty]$ by $\psi(t) := -\ln \frac{4e^t}{(1+e^t)^2}$, $t \in I$. Then for all $\varepsilon \in [0, B]$ we have

$$\frac{e^\varepsilon - 1}{2e^\varepsilon} \ln \frac{e^B + e^{2\varepsilon}}{e^B + e^\varepsilon} \leq \delta_\psi(2\varepsilon) \leq \frac{e^\varepsilon - 1}{e^\varepsilon} \ln \frac{e^B + e^{2\varepsilon}}{e^B + e^\varepsilon}.$$

Indeed, an easy calculation shows $\psi'(t) = \frac{e^t - 1}{e^t + 1}$ for all $t \in I$, and hence we obtain

$$\psi'(t) - \psi'(t-\varepsilon) = \frac{e^t - 1}{e^t + 1} - \frac{e^t - e^\varepsilon}{e^t + e^\varepsilon} = \frac{2e^t(e^\varepsilon - 1)}{(e^t + 1)(e^t + e^\varepsilon)}$$

for all $t \in [0, B]$, $\varepsilon \in [0, B]$. Consequently we have

$$\frac{e^\varepsilon - 1}{e^t + e^\varepsilon} \leq \psi'(t) - \psi'(t-\varepsilon) \leq 2 \frac{e^\varepsilon - 1}{e^t + e^\varepsilon}$$

for all $t \in [0, B]$, $\varepsilon \in [0, B]$. Furthermore, for $\varepsilon > 0$ an easy calculation gives

$$\inf_{x \in [0, B-\varepsilon]} \int_x^{x+\varepsilon} \frac{e^\varepsilon - 1}{e^t + e^\varepsilon} dt = \int_{B-\varepsilon}^B \frac{e^\varepsilon - 1}{e^t + e^\varepsilon} dt = \frac{e^\varepsilon - 1}{e^\varepsilon} \left(t - \ln(e^t + e^\varepsilon) \right) \Big|_{t=B-\varepsilon}^B = \frac{e^\varepsilon - 1}{e^\varepsilon} \ln \frac{e^B + e^{2\varepsilon}}{e^B + e^\varepsilon}.$$

From this we easily find the assertion.

The following two lemmas establish important properties of the Fenchel-Legendre bi-conjugate operation ** . We begin with

Lemma A.6 *Let $B > 0$ and $\delta : [0, B] \rightarrow [0, \infty)$ be a monotone increasing function with $\delta(0) = 0$ and $\delta(\varepsilon) > 0$ for all $\varepsilon \in (0, B]$. Then for all $\varepsilon \in (0, B]$ we have*

$$\delta^{**}(\varepsilon) > 0.$$

Proof: Let us assume that there exists an $0 < \varepsilon \leq B$ with $\delta^{**}(\varepsilon) = 0$. Then we have $(\varepsilon, 0) \in \text{Epi } \delta^{**} = \overline{\text{co Epi } \delta}$, and hence there exists a sequence $(\varepsilon_n, y_n) \in \text{co Epi } \delta$ with $\varepsilon_n \rightarrow \varepsilon$ and $y_n \rightarrow 0$. Furthermore, we have $\text{co Epi } \delta \subset \mathbb{R}^2$, and hence Carathéodory's theorem (see e.g. [23][p. 55]) guarantees that for all $n \geq 1$ there exist $\varepsilon_{n,1}, \varepsilon_{n,2}, \varepsilon_{n,3} \in [0, B]$, $y_{n,1}, y_{n,2}, y_{n,3} \in [0, \infty)$, and $\alpha_{n,1}, \alpha_{n,2}, \alpha_{n,3} \in [0, 1]$ with

$$\begin{aligned} \varepsilon_n &= \alpha_{n,1}\varepsilon_{n,1} + \alpha_{n,2}\varepsilon_{n,2} + \alpha_{n,3}\varepsilon_{n,3}, \\ y_n &= \alpha_{n,1}y_{n,1} + \alpha_{n,2}y_{n,2} + \alpha_{n,3}y_{n,3}, \\ 1 &= \alpha_{n,1} + \alpha_{n,2} + \alpha_{n,3}, \\ y_{n,i} &\geq \delta(\varepsilon_{n,i}), \end{aligned} \quad i = 1, \dots, 3.$$

In addition we may assume $\varepsilon_{n,1} \leq \varepsilon_{n,2} \leq \varepsilon_{n,3}$ without loss of generality. Since this yields $\varepsilon_n = \alpha_{n,1}\varepsilon_{n,1} + \alpha_{n,2}\varepsilon_{n,2} + \alpha_{n,3}\varepsilon_{n,3} \leq \varepsilon_{n,3}$ we find $y_{n,3} \geq \delta(\varepsilon_{n,3}) \geq \delta(\varepsilon_n) \geq \delta(\frac{\varepsilon}{2}) > 0$ for large n . Recalling $y_n \rightarrow 0$ we thus obtain $\alpha_{n,3} \rightarrow 0$ which implies both $\alpha_{n,1} + \alpha_{n,2} \rightarrow 1$ and $\alpha_{n,1}\varepsilon_{n,1} + \alpha_{n,2}\varepsilon_{n,2} \rightarrow \varepsilon$. However, the latter convergence gives $\frac{\varepsilon}{2} \leq \alpha_{n,1}\varepsilon_{n,1} + \alpha_{n,2}\varepsilon_{n,2} \leq (\alpha_{n,1} + \alpha_{n,2})\varepsilon_{n,2}$ for large n , and hence we have $\varepsilon_{n,2} \geq \frac{\varepsilon}{4}$ for large n . Again this shows $y_{n,2} \geq \delta(\varepsilon_{n,2}) \geq \delta(\frac{\varepsilon}{4}) > 0$ for large n , and thus we find $\alpha_{n,2} \rightarrow 0$. Obviously, this yields both $\alpha_{n,1} \rightarrow 1$ and $\alpha_{n,1}\varepsilon_{n,1} \rightarrow \varepsilon$, and hence we obtain $\varepsilon_{n,1} \geq \frac{\varepsilon}{4}$ for large n . Finally, this gives $y_{n,1} \geq \delta(\varepsilon_{n,1}) \geq \delta(\frac{\varepsilon}{4}) > 0$ for large n and therefore we find $\alpha_{n,1} \rightarrow 0$ which contradicts the already found convergence $\alpha_{n,1} \rightarrow 1$. ■

Lemma A.7 *Let $B > 0$, and $\delta : [0, B] \rightarrow [0, \infty)$ be a continuous function with $\delta(0) = 0$. We define $\tilde{\delta} : [0, B] \rightarrow [0, \infty)$ by $\tilde{\delta}(\varepsilon) := \inf_{\varepsilon' \geq \varepsilon} \delta(\varepsilon')$, $\varepsilon \in [0, B]$. Then $\tilde{\delta}$ is monotonously increasing, and for all $\varepsilon \in [0, B]$ we have*

$$\delta^{**}(\varepsilon) = \tilde{\delta}^{**}(\varepsilon).$$

In addition, if $\delta(\varepsilon) > 0$ for all $\varepsilon \in (0, B]$, then $\tilde{\delta}(\varepsilon) > 0$ for all $\varepsilon \in (0, B]$.

Proof: The first assertion is trivial and the third assertion directly follows from the continuity of δ . Therefore, it remains to show

$$\text{co Epi } \delta = \text{co Epi } \tilde{\delta} \tag{62}$$

since this equation immediately yields $\delta^{**} = \tilde{\delta}^{**}$. To establish (62) we first observe that $\tilde{\delta}(\varepsilon) \leq \delta(\varepsilon)$ for all $\varepsilon \in [0, B]$ and hence we have $\text{co Epi } \delta \subset \text{co Epi } \tilde{\delta}$. To prove the converse inclusion observe that it suffices to show $(\varepsilon, \tilde{\delta}(\varepsilon)) \in \text{co Epi } \delta$ for all $\varepsilon \in [0, B]$. Furthermore, we have $\tilde{\delta}(0) = 0 = \delta(0)$ and $\tilde{\delta}(B) = \delta(B)$ and hence we can restrict our considerations to pairs $(\varepsilon, \delta(\varepsilon))$ for $\varepsilon \in (0, B)$. Therefore let us fix an $\varepsilon \in (0, B)$. By the definition of $\tilde{\delta}$ we then find an $\varepsilon_+ \in [\varepsilon, B]$ with $\delta(\varepsilon_+) = \tilde{\delta}(\varepsilon)$. Furthermore, we have $\delta(0) \leq \tilde{\delta}(\varepsilon) \leq \delta(\varepsilon)$ and hence the intermediate value theorem applied to the continuous function δ gives us an $\varepsilon_- \in [0, \varepsilon]$ with $\delta(\varepsilon_-) = \tilde{\delta}(\varepsilon)$. Now, there exists an $\alpha \in [0, 1]$ with $\varepsilon = \alpha\varepsilon_+ + (1 - \alpha)\varepsilon_-$ and since our previous considerations showed

$$\tilde{\delta}(\varepsilon) = \alpha\tilde{\delta}(\varepsilon) + (1 - \alpha)\tilde{\delta}(\varepsilon) = \alpha\delta(\varepsilon_+) + (1 - \alpha)\delta(\varepsilon_-)$$

we obtain $(\varepsilon, \tilde{\delta}(\varepsilon)) \in \text{co Epi } \delta$. ■

References

- [1] P.L. Bartlett, M.I. Jordan, and J.D. McAuliffe. Convexity, classification, and risk bounds. *J.A.S.A.*, to appear, to appear, 2006. <http://stat-www.berkeley.edu/tech-reports/638.pdf>.
- [2] F.A. Behringer. Convexity is equivalent to midpoint convexity combined with strict quasiconvexity. *Optimization*, 24:219–228, 1992.
- [3] C. Bennett and R. Sharpley. *Interpolation of Operators*. Academic Press, Boston, 1988.
- [4] G. Blanchard, O. Bousquet, and P. Massart. Statistical performance of support vector machines. *Ann. Statist.*, submitted, 2004.
- [5] D. Butnariu and A.N. Iusem. *Totally Convex Functions for Fixed Points Computation and Infinite Dimensional Optimization*. Kluwer, Dordrecht, 2000.
- [6] A. Caponnetto. A note on the role of squared loss in regression. Technical report, Massachusetts Institute of Technology, 2005. <http://cbcl.mit.edu/projects/cbcl/publications/ps/caponnetto-squareloss-6-05.pdf>.
- [7] C. Castaing and M. Valadier. *Convex Analysis and Measurable Multifunctions*. Springer-Verlag, Berlin, 1977.
- [8] A. Christmann and I. Steinwart. Consistency and robustness of kernel based regression. *Bernoulli*, submitted, 2005. <http://www.c3.lanl.gov/~ingo/publications/ann-04b.pdf>.
- [9] O. Dekel, S. Shalev-Shwartz, and Y. Singer. Smooth ε -insensitive regression by loss symmetrization. *J. Mach. Learn. Res.*, 6:711–741, 2005.
- [10] L.M. Graves. *The Theory of Functions of Real Variables*. McGraw-Hill, New York, 1956.
- [11] J.A. Hartigan. Estimation of a convex density contour in 2 dimensions. *J. Amer. Statist. Assoc.*, 82:267–270, 1987.
- [12] T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning*. Springer, New York, 2001.
- [13] K-U. Höffgen and Hans-U. Simon. Robust trainability of single neurons. In *Proceedings of the Computational Learning Theory (COLT) Conference*, pages 428–438, 1992.
- [14] P.J. Huber. Robust estimation of a location parameter. *Ann. Math. Statist.*, 35:73–101, 1964.
- [15] H. Kestelman. *Modern Theories of Integration*. Dover, New York, 1960.
- [16] E.S. Levitin and B.T. Polyak. Convergence of minimizing sequences in conditional extremum problems. *Soviet Math. Dokl.*, 7:764–767, 1966.
- [17] Y. Lin. Support vector machines and the Bayes rule in classification. *Data Mining and Knowledge Discovery*, 6:259–275, 2002.
- [18] Y. Lin. A note on margin-based loss functions in classification. *Statist. Probab. Lett.*, 68:73–82, 2004.

- [19] Y. Lin, Y. Lee, and G. Wahba. Support vector machines for classification in nonstandard situations. *Machine Learning*, 46:191–202, 2002.
- [20] D.W. Müller and G. Sawitzki. Excess mass estimates and tests for multimodality. *J. Amer. Statist. Assoc.*, 86:738–746, 1991.
- [21] W. Polonik. Measuring mass concentrations and estimating density contour clusters—an excess mass approach. *Ann. Statist.*, 23:855–881, 1995.
- [22] B.T. Polyak. Existence theorems and convergence of minimizing sequences for extremal problems with constraints. *Soviet Math. Dokl.*, 7:72–75, 1966.
- [23] R.T. Rockafellar and R.J.-B. Wets. *Variational Analysis*. Springer, 1998.
- [24] A. Smola, B. Schölkopf, and K.-R. Müller. Convex cost functions for support vector regression. In L. Niklasson, M. Bodén, and T. Ziemke, editors, *Proceedings of the International Conference on Artificial Neural Networks*, Perspectives in Neural Computing, pages 99–104, Berlin, 1998. Springer.
- [25] I. Steinwart. Sparseness of support vector machines. *J. Mach. Learn. Res.*, 4:1071–1105, 2003.
- [26] I. Steinwart. Consistency of support vector machines and other regularized kernel machines. *IEEE Trans. Inform. Theory*, 51:128–142, 2005.
- [27] I. Steinwart, D. Hush, and C. Scovel. A classification framework for anomaly detection. *J. Mach. Learn. Res.*, 6:211–232, 2005.
- [28] I. Steinwart and C. Scovel. Fast rates for support vector machines using Gaussian kernels. *Ann. Statist.*, submitted, 2004. <http://www.c3.lanl.gov/~ingo/publications/ann-04a.pdf>.
- [29] A. Tewari and P.L. Bartlett. On the consistency of multiclass classification methods. In *Proceedings of the 18th Annual Conference on Learning Theory, COLT 2005*, pages 143–157. Springer, 2005.
- [30] A.B. Tsybakov. On nonparametric estimation of density level sets. *Ann. Statist.*, 25:948–969, 1997.
- [31] V.N. Vapnik. *The Nature of Statistical Learning Theory*. Springer, New York, 1998.
- [32] V.N. Vapnik. *Statistical Learning Theory*. Wiley, 1998.
- [33] C. Zălinescu. On uniformly convex functions. *J. Math. Anal. Appl.*, 95:344–374, 1983.
- [34] C. Zălinescu. *Convex Analysis in General Vector Spaces*. World Scientific, New Jersey, 2002.
- [35] T. Zhang. Statistical analysis of some multi-category large margin classification methods. *J. Mach. Learn. Res.*, 5:1225–1251, 2004.
- [36] T. Zhang. Statistical behaviour and consistency of classification methods based on convex risk minimization. *Ann. Statist.*, 32:56–134, 2004.